6.5.2 Heavy Ion (HI) Driver Subsystems

The Prometheus-H heavy ion driver is based on space charge limited acceleration and transport of a heavy ion beam in a simple alternating magnetic quadrupole focussing channel. The channel containing the intense beam is formed by a slowly varying FODO lattice (focus-drift-defocus-drift). The Prometheus-H Heavy Ion Driver illuminates an indirect-drive target from two sides. The major subsystem elements of the heavy ion driver are:

(1) A single-beam induction linear accelerator (LINAC) produces a train of identical current pulses running in a high repetition rate burst mode.

(2) Storage rings accumulate the LINAC current pulses and act as a series of delay lines. Beams generated over microseconds are conditioned to reach the target within hundreds of picoseconds of each other. Storage rings are located where the beams segregate into main and precursor beams. Some of the beamlets become the two precursor beams, which contain roughly 30% of the beam energy and have a duration 4 to 5 times that of the main current pulses. The combination of the single beam LINAC and the storage rings was selected in Prometheus-H as opposed to a multiple beam accelerator usually chosen in earlier studies.

(3) Combination bunching/drift sections are downstream of the LINAC and storage rings. Their function is to compress the beams longitudinally. The precursor and main beams share a common tunnel, but they have different requirements and are handled separately in these subsystems. Following the buncher/drift section, all beamlets are divided into two identical sets and directed toward opposite sides of the target chamber.

(4) Final focus sections on each side of the reactor chamber compress the beamlets radially, coalesces the main beamlets into a main beam, and merges this beam with the end of the precursor beam just inside the target chamber. The merged beams then propagate across the target chamber to opposite sides of the indirect-drive target. The entire Prometheus-H heavy ion driver system is cycled at 3.54 Hz.

Prometheus-H heavy ion driver is composed of subsystems and individual modules which are subject to parallel optimization during the years 1992-2030. Each of these subsystems is described below. A general schematic illustrating the geometrical relationship of the main single beam LINAC to the storage rings, bunchers, final focus, and target chamber is provided below in Figure 6.5.2-1.

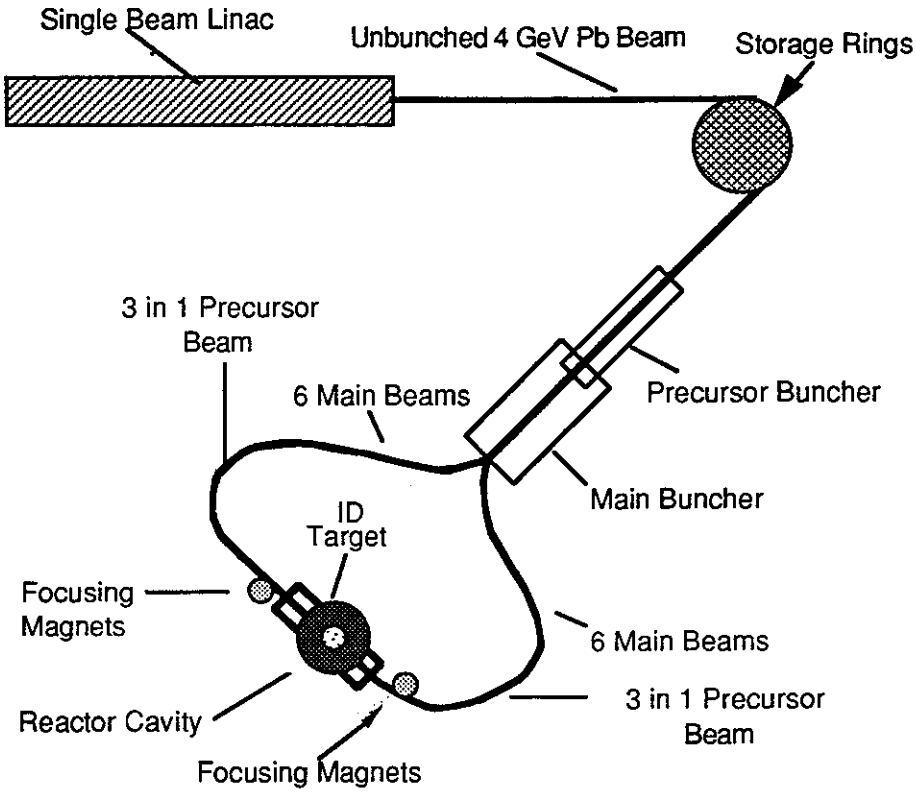

Figure 6.5.2-1 The Prometheus-H Heavy Ion Driver consists of a single beam LINAC (consisting of a high brightness injector, ramp gradient accelerator, and main accelerator) operated in a burst mode to generate 18 beamlets, 12 storage rings, focusing magnets, and two channel transport for a 2-sided indirect drive DT target.

To meet the overall IFE reactor requirements, the overall heavy ion driver is cycled at 3.54 Hz. This driver is based on space charge limited acceleration and transport of a heavy ion beam in a simple alternating magnetic quadrupole focussing channel. As noted previously, the channel containing the intense HI beam is formed by a slowly varying FODO lattice, whose parameters are illustrated in Figure 6.5.2-2.

Figure 6.5.2-2 defines the quadrupole spacing, L, the quadrupole length, $l_q = \eta L$, the space between the quadrupoles, $L(1 - \eta)$, the beam envelope of radius a, the heavy ion beam trajectory, the alternating focus and defocus (FODO) magnets, the betatron wavelength, λ_b, and the phase advance, $\sigma = 2\pi \ 2L/\lambda_b$.

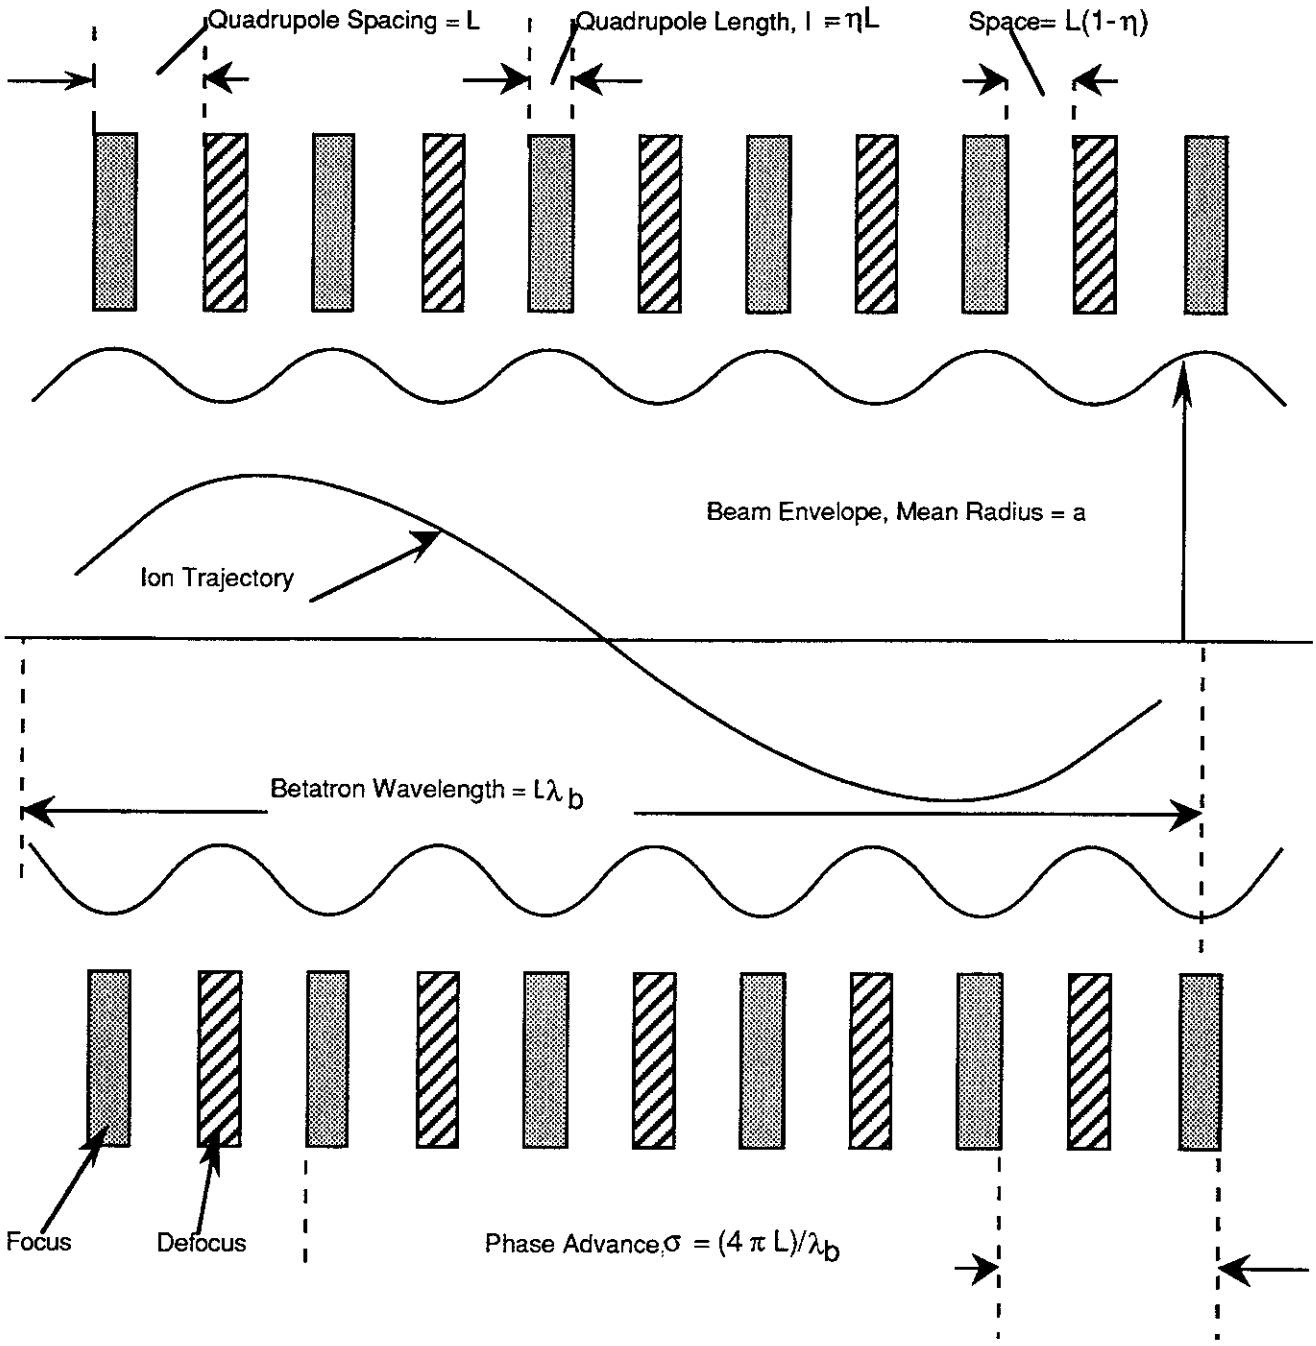

Figure 6.5.2-2 Schematic of the Prometheus-H Quadrupole Focussing Channel

The HI beam envelope has a periodicity equal to that of the quadrupoles, but the trajectory of an individual ion within the channel is characterized by a betatron wavelength. The betatron wavelength, λ_b, is indicative of the balance between the radial space charge forces on the beam and restoring force of the quadrupole channel. The betatron period divided by four, $\lambda_b/4$, is the distance in which the HI

beam would reach its maximum displacement after receiving an angular kick. An infinite betatron period (or zero phase advance) corresponds to having no net restoring force on the ion.

As indicated in Figure 6.5.2-3, at any location along the beamline, the trajectories of the individual ions fall within an envelope in phase space (x-x'). The area of the envelope is related to $\pi\varepsilon$, where ε is the beam emittance.

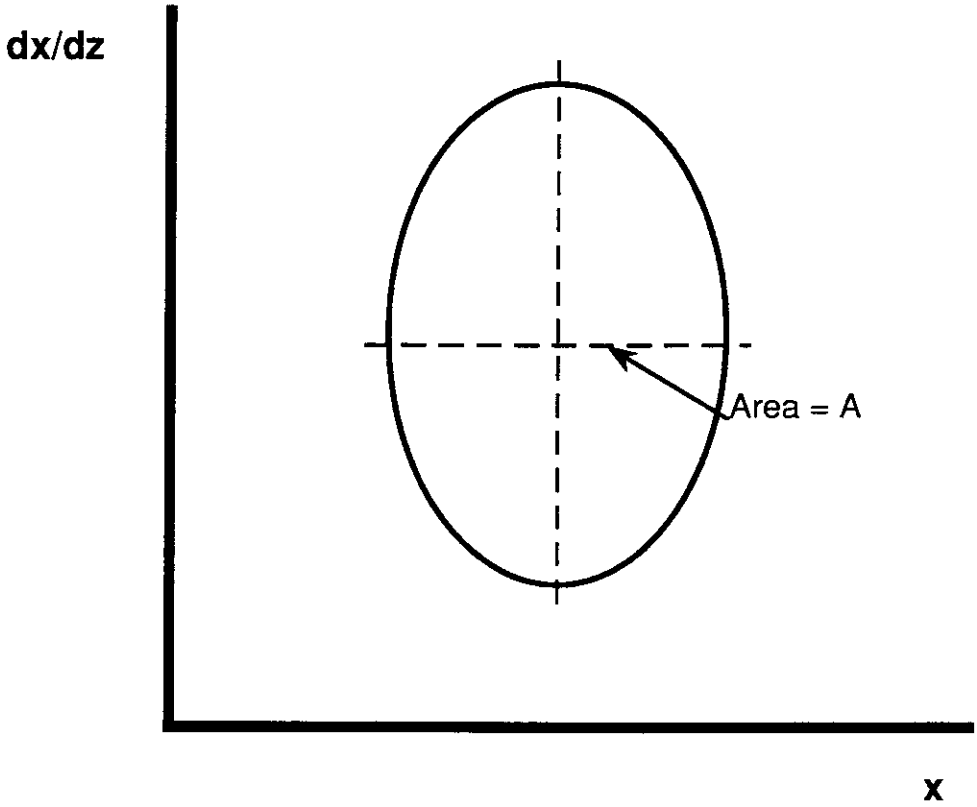

Figure 6.5.2-3 Heavy Ion Trajectories Lie Within this Elliptical Envelope in Phase Space.

For the coasting HI beam, the area is constant, although the aspect ratio changes. As the HI beam accelerates, the area decreases. The normalized emittance defined by the expression:

$$\varepsilon_N = \beta\gamma\varepsilon \qquad (6.5.2\text{-}1)$$

The strength of the FODO channel focussing depends not only on the physical spacing of the magnets but the quadrupole gradient as well; the magnetic parameters are illustrated in Figure 6.5.2-4. (Note: For the type of magnet proposed, "poletip" is a misnomer. The poletip field is defined as the magnetic field at the beam edge.)

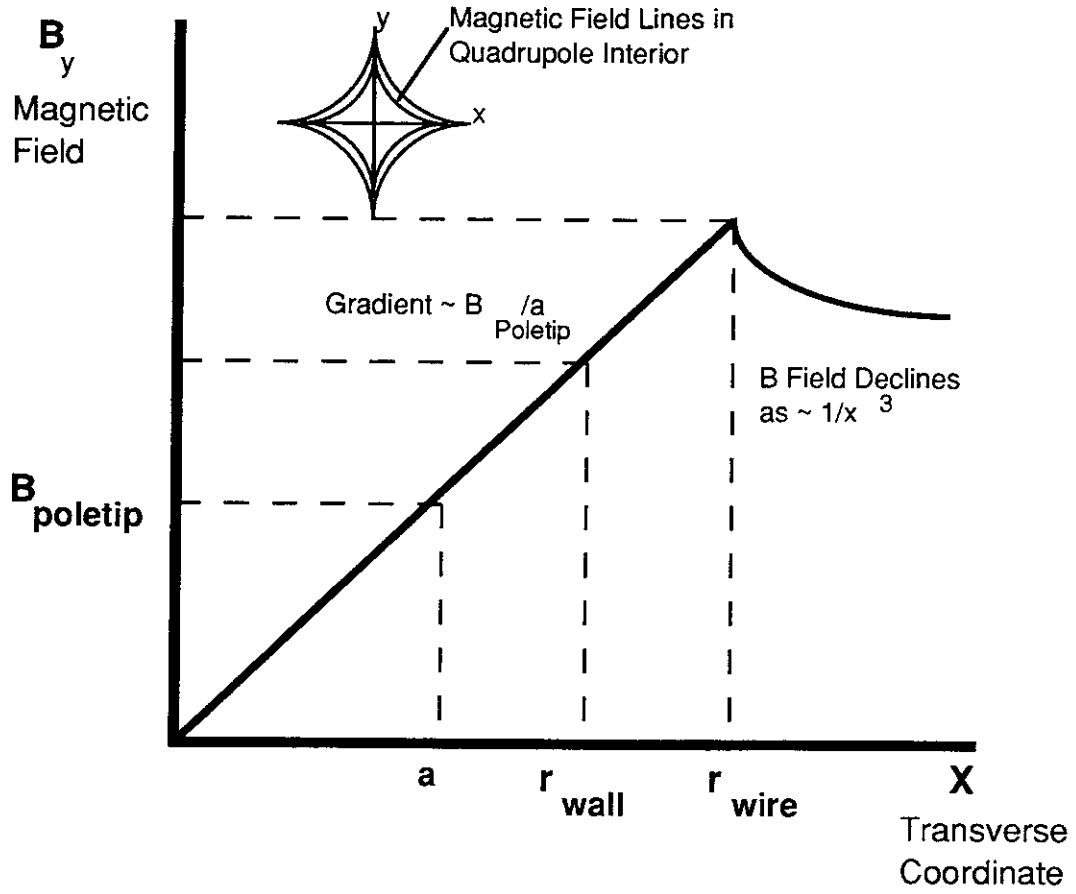

Figure 6.5.2-4 Relationships Between Magnetic Fields at Poletips, Walls, and Wires

The lattice parameters vary along the driver according to the local beam energy and/or current; the acceleration and bunching schedules (see Figures 6.5.2-5 and 6.5.2-6) which determine V(z) and I(z) are free parameters.

Note that the ramp portion of the LINAC commences the prebunching of the HI beams. Beyond the ramp gradient region of the LINAC, bunching continues.

Illustrated in Figure 6.5.2-6 is the heavy ion beam current as a function of time at a given location in the Prometheus-H heavy ion driver. In this case, the HI pulse duration, $\tau = \delta/v$ (where δ is the HI pulse length in m, and v is the bunch velocity in m/s).

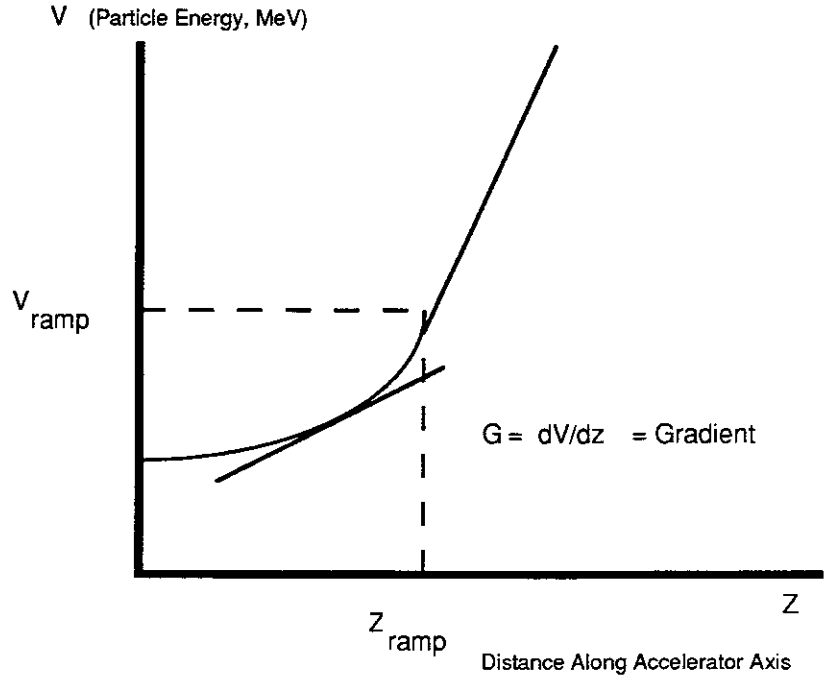

Figure 6.5.2-5 Relationship Between Particle Energy, V, and Accelerator Gradient, G, in LINAC

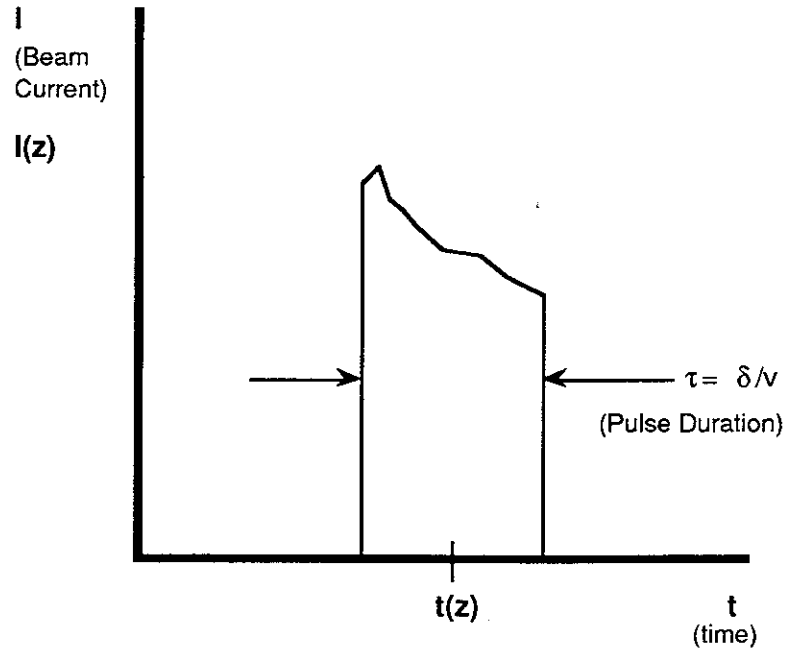

Figure 6.5.2-6 Heavy Ion beam current as a function of time at a given location in the Prometheus Heavy Ion Driver.

Simple equations for the space charge-limited current, the emittance, and the quadrupole focussing strength were used to derive three scaling parameters. These characterize the driver and lattice variables throughout the bunch acceleration and compression processes in terms of the local beam energy and peak current. Practical limitations on physical parameters (e.g., component access, currents in magnet windings, minimum depressed tune, etc.) were used to constrain the parameters to realizable configurations of the driver. Given a set of basic beam parameters at the target (total beam energy, ion kinetic energy, number of beamlets, pulse length, etc.) the cost of electricity from the plant minimizes at specific values of the three parameters; different values are obtained if the LINAC and buncher costs are considered separately. The scaling laws, practical constraints, and driver variables in terms of the three scaling parameters are given below.

Equations Governing Transport in FODO - To define the space-charge limited current, I_{max}, the following expressions are valid:

$$I_{max} = 1.56\times10^{7}\,\sigma_o^2\left(\frac{a}{2L}\right)^2\frac{A}{Z}(\beta\gamma)^3 \qquad (6.5.2\text{-}2)$$

where $\sigma = 2\pi/\lambda_o$ is the depressed tune (and is λ_o the wavelength of the ion trajectory, a is the beam radius, L is the distance between alternating quadrupoles, A is the atomic weight (which in the case of Pb requires A = 208), Z = 2 is the ion charge state, $\beta = v_i/c$ (where v_i = ion velocity), similarly $\gamma = 1/(1-\beta^2)^{1/2} \sim 1 + \beta^2/2 \sim 1.0029$ for these non-relativistic ions. To define the superconducting quadrupole limitation, the following expression for the poletip magnetic field, $B_{poletip}$, is:

$$B_{poletip} = 3.13\frac{A}{Z}\left(\frac{a}{L^2}\right)(\beta\gamma)\left(\frac{1}{\eta}\right) \qquad (6.5.2\text{-}3)$$

where the constant ($\xi = 3.13$) is based upon recent experience with Superconducting Supercollider (SCC) quadrupoles, η = quadrupole packing fraction = $1_q/L$ (1_q = quadrupole length), and $B_{poletip}$ is the magnetic field at the beam edge. In order to calculate the constant normalized emittance, ε_N, the following equations were used:

$$\varepsilon_N = \frac{2\pi\, a^2\, \beta\, \gamma}{\lambda_\beta}$$

$$\varepsilon_N = \frac{\sigma\, a^2\, \beta\, \gamma}{2\, L}$$

$$\beta\gamma = \frac{v}{c\sqrt{1 - \frac{v^2}{c^2}}}$$

(6.5.2-4)

A summary of the heavy ion driver parameters is provided below in Table 6.5.2-1.

Table 6.5.2-1
Summary of Heavy Ion Driver Parameters

#	Parameter	Symbol	Variable Voltage	Constant Voltage
1	Depressed Tune	σ	V^{κ}	$I^{-\zeta}$
2	Current	I	$V^{1-\alpha}$	I
3	Beam Radius	a	$V^{\alpha/2-\kappa-1/4}$	$I^{\zeta-1/2}$
4	Pulse Duration	τ	$V^{\alpha-1}$	I^{-1}
5	Bunch Length	δ	$V^{\alpha-1/2}$	I^{-1}
6	Quad Spacing	L	$V^{\alpha-\kappa}$	$I^{\zeta-1}$
7	Packing Fraction	η	$V^{-3/2\alpha+\kappa+1/4}$	$I^{3/2-\zeta}$
8	Quad Length	l_q	$V^{1/4-\alpha/2}$	$I^{1/2}$
9	Ion Charge	Q		

Additional constraints include: For adequate access to the HI beamlines, $\eta < 0.8$ everywhere. The control of the final focus magnetic lens aberration requires that the ratio of the beam radius to the quad length satisfy the following inequality:

$$a/l_q < 0.25 \qquad (6.5.2\text{-}5)$$

throughout the HI driver FODO system. For the superconducting magnets, it is necessary that the current density be kept below its local limit; SSC-type conductors are assumed and the limits are never approached in Prometheus-H.

The Prometheus-H driver subsystems generally stay within the bounds of demonstrated or readily extrapolated technologies. There are some exceptions, which are listed here and will be discussed in more detail in the following sections. First, the entire Prometheus-H driver (LINAC, storage ring, and buncher sections) is designed around a very aggressive transport schedule. The current is maintained at the local

space charge limit of a magnetic quadrupole focussing channel throughout the driver, and the depressed tune is allowed to drop as the beam approaches the target. The ability to transport the current at its space charge limit over long distances is identified as a Critical Issue for the Heavy Ion Driver since it heavily impacts the driver cost.

Second, the ability to maintain the high current heavy ion beams in the storage rings for millisecond time scales is also identified as a Critical Issue, since the use of a cost-effective Single Beam LINAC for this application hinges on the ability to delay the arrival time of successively generated bunches. Last, the final stage of the transport uses an admittedly speculative process, self-pinched transport in a narrow channel; this is not listed as a Critical Issue because there are acceptable (albeit less attractive) workarounds if the process proves to be unrealizable in practice.

6.5.2.1 Single Beam LINAC - The high repetition rate burst mode Single Beam LINAC comprises two portions, the injector and the main LINAC. The injector includes the ion source and whatever preacceleration and beam manipulation is required to produce a suitable beam for the main LINAC. The two sections of the main LINAC, the ramped gradient and the fixed gradient sections, differ in the allowed voltage gradient as a function of position, but are otherwise identical in function: they both accelerate a single space-charge-limited beam at the maximum allowed rate. The main components of the single beam LINAC are the induction modules that accelerate the beam and the superconducting quadrupole magnets that form the transport channel.

6.5.2.1.1 Injector - The baseline Prometheus-H driver design calls for the injector to deliver a beam with specific properties at the injection point. The Prometheus-H heavy ion driver injector parameters are summarized below in Table 6.5.2-2.

Table 6.5.2-2
Summary of Injector Parameters

#	Parameter	Symbol	Choice	Unit
1	Ion Charge State	q	Pb^{+2}	
2	Energy	V	6 .0	MeV
3	Beam Current	I	14	A
4	Pulse Duration	τ	15.5	μs
5	Pulse Shape		Nominally Square	
6	Beam Radius	a	9.4	cm
7	Current Density	$I/\pi a^2$	50	mA/cm^2
8	Normalized Emittance	ε	< .00001	m-rad
9	Pulses in Burst Mode		18	
10	Pulse Repetition Rate		30	kHz
11	Time Between Bursts		200	ms

In defining the injector requirements it is anticipated that the techniques in use today will be extrapolated to meet these requirements. The problems in the injector are the production of the total current, with the required brightness, at a sufficiently high repetition rate and duty factor, without generating an unacceptable gas load for the rest of the LINAC, and to do all this reliably. The development of the injector is identified as one of the key issues for a heavy ion driver.

The beam was assumed to be space charge limited at the injector exit, rather than source limited. Although the current is high, the cross section of the beam is large and the required electrical current density is below 60mA/cm^2. Nearly 600 mA/cm^2 has been demonstrated using high brightness proton sources and the C^{1+} source developed for the ILSE program delivers 30mA/cm^{-2} at 2MeV, but when adjusted for the charge and mass (the space charge-limited current density scales as $(Q/A)^{-1/2}$, where Q is the ion charge and A is the atomic number) the Pb^{2+} beam of Prometheus-H is seen to have a stricter requirement. It is possible that multiple beamlets can be generated and combined before injection into the LINAC, since the brightness of these beams exceeds the driver requirement; excess brightness provides a margin against phase space dilution downstream as the individual beamlets are combined. It would be preferable if the beam could be generated in a single stage.

Source techniques in use today will probably be used in the driver injector. The ions could be extracted from a pulsed lead plasma of number density around 10^{13} cm^{-3}; low ion temperature rf discharge plasmas can approach those densities over large volumes. Plasma control via a pulsed extraction electrode, currently used with the C^+ source mentioned above, can switch the flow of ions from a multi-aperture source on and off. It is possible that the ions generated in the rf discharge plasma will be created in a lower charge state (Pb^{1+}) or in a mixture of states. This would require charge state selection; if the 1+ state is dominant, the beam would require 6 MV of acceleration to reach the ion injection energy, followed by stripping to the 2+ state.

The high voltage techniques used in the driver injector will also probably be similar to today's, with adjustments made for the high duty factor. Burst repetition rate in the driver system is limited by the source duty factor at the high end, and practically by the beam storage time at the other extreme (and ultimately by the 5 Hz target interaction rate). The ~15 μs pulse duration at the injector is sufficiently long that many of the issues of steady state operation must be faced: gas loading, voltage standoff, etc. The repetition rate set by the storage rings is such that during a ~2 ms burst the duty factor of the injector is running at around 10%.

One option for the voltage source for the injector is a series of induction modules ganged together and applying a large positive potential to the source, which extends through the modules on a stalk. This has an advantage in that it shares a common high voltage technology with the rest of the LINAC. It is not clear that this is desirable, since the core area that is required to provide the 45 V-s for a 3 MV, 15 μs pulse is nearly 30 m^2 for a 1.5 T flux swing. Although core volume is minimized by using long, small diameter cylinders, even a 30 m long cylinder has a minimum volume of over 90 m^3. Even at a low value of 100 J/m^3 loss per pulse in the Metglas, this is nearly a megawatt in heat just in the core. In contrast, the beam power is only 60 kW.

Another option might be a dc high voltage power supply. The ~15 ms pulse duration at the injector is sufficiently long that many of the issues of steady state operation must already be faced: gas loading, voltage standoff, etc. Although the peak current is 14 A, the true average beam current is only 20 mA and the average current during the burst is only 1.5 A. DC power supplies capable of 3 MV at tens of milliamperes are available commercially now, and the design is such that 2 A of average current should be possible.

Source techniques in use today will probably be adapted for use in the driver injector. The ions could be extracted from a pulsed lead plasma of number density around 10^{13} cm^{-3}; low ion temperature rf discharge plasmas can approach those densities over large volumes. The problem of gas loading can be mitigated to some extent by puffing the gas in before each shot; unlike hydrogen, lead is readily pumped. Plasma control via a pulsed extraction electrode, currently used with the carbon arc source mentioned above, can switch the flow of ions from a multi-aperture source on and off.

It is possible that the ions generated in the rf discharge plasma will be created in a lower charge state (Pb^{1+}) or in a mixture of states. This would require charge state selection; if the 1+ state is dominant, the beam would require 6 MV of acceleration to reach the ion injection energy, followed by stripping to the 2+ state for propagation in the rest of the system.

A plot of the acceleration gradient, dV/dz, in the ramp-gradient portion of the Prometheus-H driver LINAC is shown in Figure 6.5.2-7.

The variation of beam energy with distance along the accelerator axis, V(z), for a displacement of up to 500 m is plotted in Figure 6.5.2-8,.

The variation of heavy ion beam energy with distance through the constant gradient portion of the Prometheus-H accelerator is illustrated in Figure 6.5.2-9.

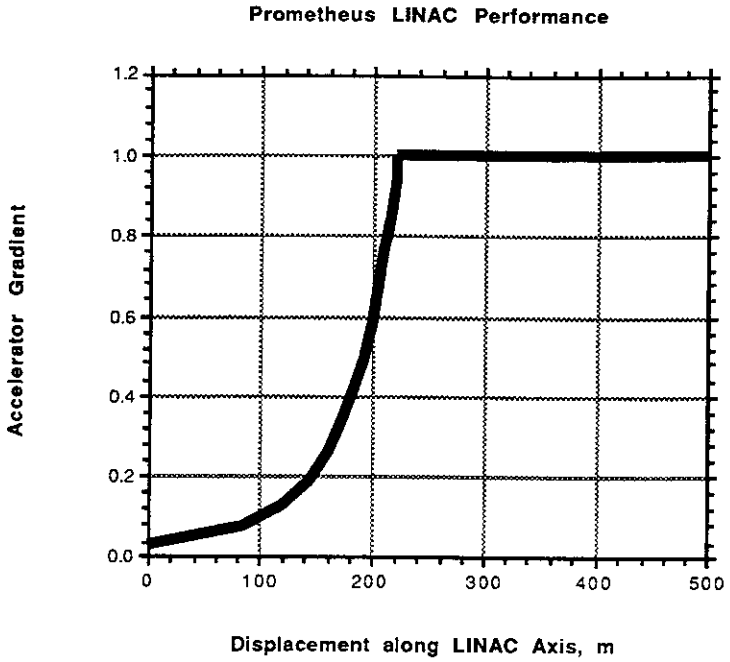

Figure 6.5.2-7 Dependence of accelerator gradient on the displacement along the LINAC axis, z, in the ramp-gradient section of the Prometheus-H heavy ion driver.

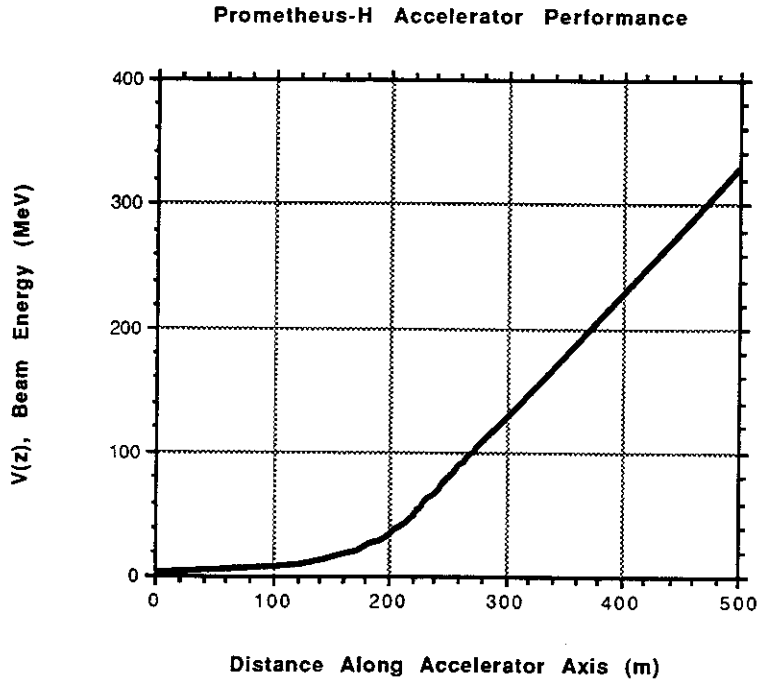

Figure 6.5.2-8 Variation of heavy ion beam energy with distance along accelerator axis in ramp-gradient portion of accelerator.

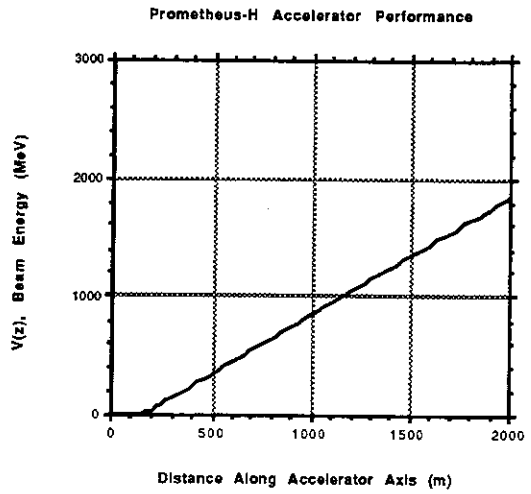

Figure 6.5.2-9 Variation of beam energy with displacement along axis for both ramped gradient and constant gradient accelerators.

Main LINAC - The heavy ion beam is accelerated at the maximum possible rate in the LINAC. In the ramped gradient section, the local voltage gradient is limited to a value proportional to the local ion energy. This is necessary so that the velocity gained at different portions of the beam as it traverses a given gap are sufficient to bunch the beam without grossly affecting the energy gained by the pulse as a whole. The end of the ramped gradient section is defined to be point where the maximum gradient is reached, at the voltage V_{ramp}. The fixed gradient section is just that, and the energy gain is linear with distance from V_{ramp} to the end of the LINAC. The variation in velocity gained as the beam traverses a gap is sufficiently low that the beam can be compressed without limiting the gradient.

The maximum gradient is taken to 1 MV/m, which should be a conservative value for a pulsed system. Graded DC accelerator columns typically run at 22 kV/cm in multi-electrode systems (e.g. Van de Graaf accelerators), and as high as 35 kV/cm in compact few-gap systems such as guns. The limitation is usually insulator flash-over. Presumably a pulsed gradient will perform at least as well.

Based on this choice for maximum gradient, The driver scaling parameters in the LINAC were found to be $\alpha = 0.2$, $\kappa = -0.15$; the resulting acceleration schedule for Prometheus-H is shown in Figure 6.5.2-9. The various accelerator and beam parameters associated with these scaling parameters are shown as a function of distance along the LINAC in the subsequent Figures 6.5.2-10 through 15.

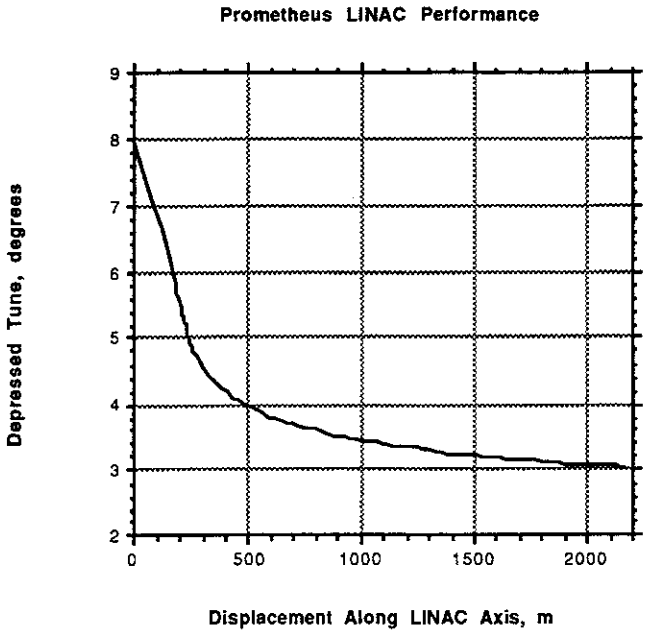

Figure 6.5.2-10 Depressed tune, σ, vs. displacement along the LINAC axis

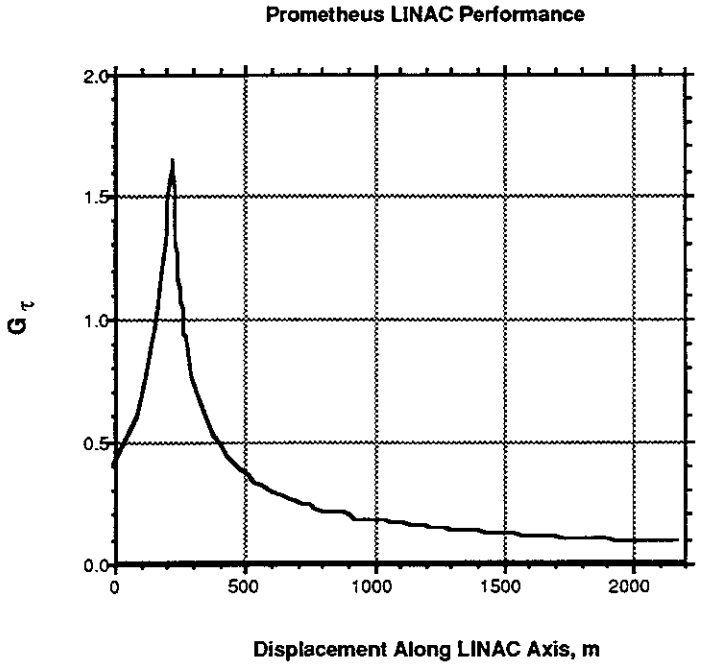

Figure 6.5.2-11 G_τ as a function of displacement along LINAC axis.

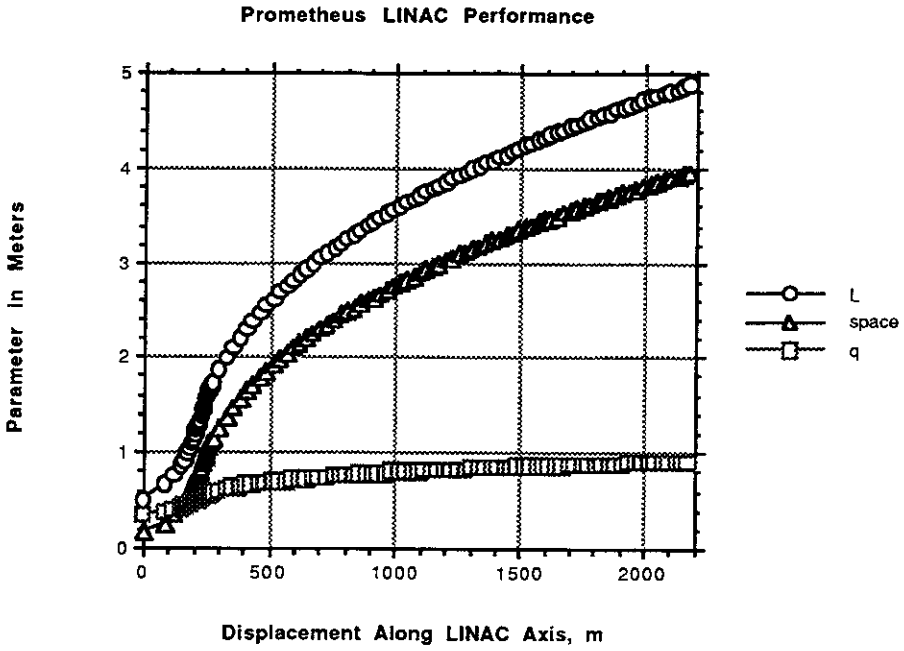

Figure 6.5.2-12 Variations of Displacement between Quads (L), blank space between Quads (space) and Quadrupole Length (q) as functions of displacement along accelerator axis.

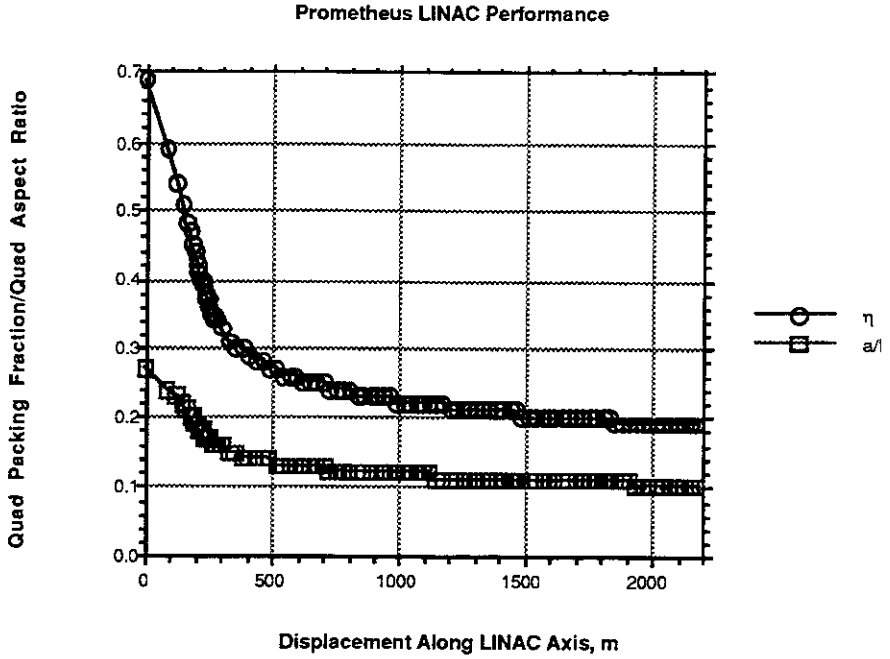

Figure 6.5.2-13 Variation of Quad Packing Fraction (η) and Quad Aspect Ratio (a/l) as functions of displacement along LINAC axis (z).

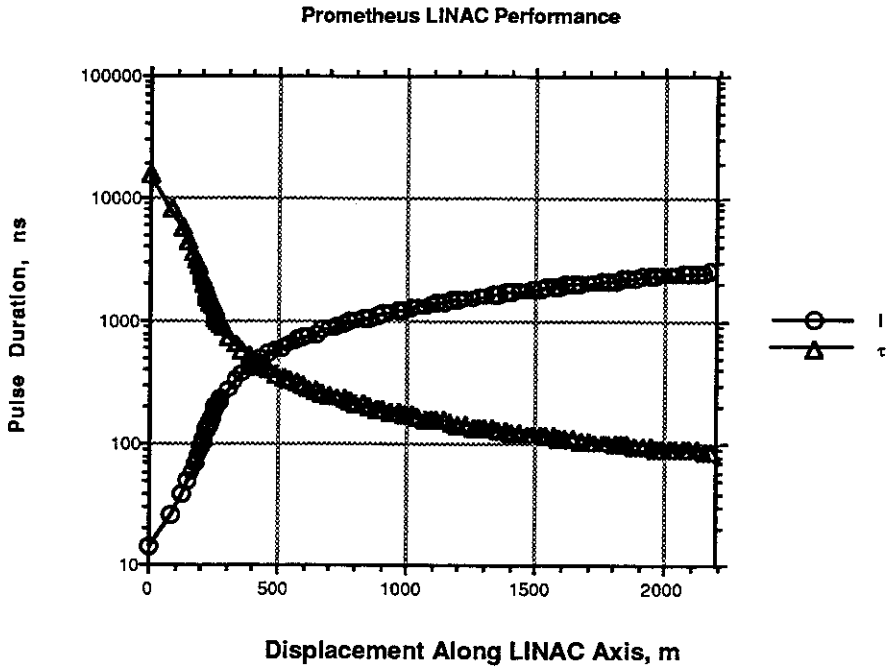

Figure 6.5.2-14 Dependence of HI beam current (I) and pulse duration (τ) on location along LINAC axis (z).

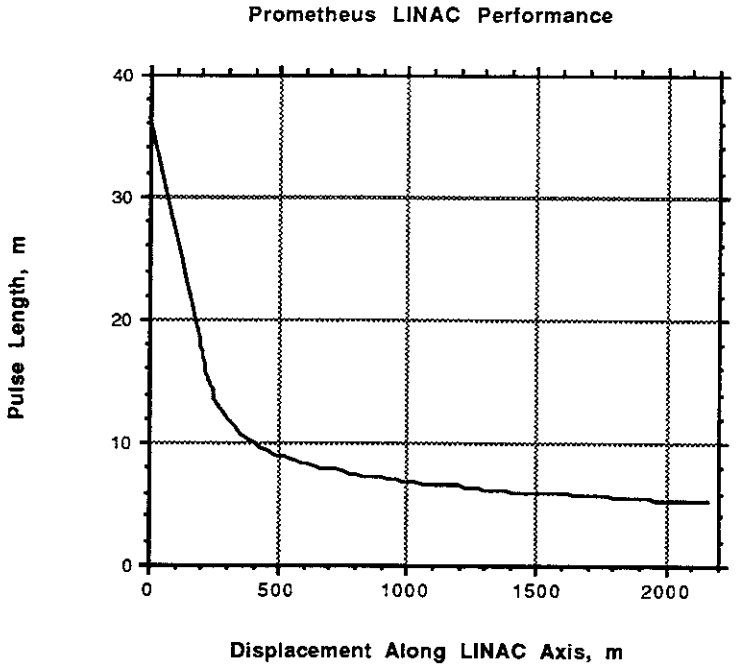

Figure 6.5.2-15 Pulse length (τ) of HI beam as a function of displacement along LINAC axis (z).

The depressed tune (betatron phase advance per quadrupole pair) decreases from 8 degrees to 3; most of the drop occurs in the first 225 m, i.e., the ramped gradient section. Displayed in Figure 6.5.2-11 is the volt-seconds per meter required to maintain the induction modules; it is a product of the gradient and the pulse duration. This number, when divided by the nominal flux swing in the cores, yields the minimum difference in inner and outer core radii; the actual difference is larger to account for core packing fraction, i.e. the average density of Metglas in the core windings when insulation and other factors are included. The quadrupole center-to-center spacing L, the length lq, and the space available for access between the magnets are shown in Figure 6.5.2-12.

In Figure 6.5.2-13, the quadrupole longitudinal packing fraction η, the ratio of quadrupole length to center-to-center spacing, and a/l_q, the ratio of beam radius to quadrupole length as functions of the distance along the accelerator axis (m). (The beam radius is a constant value of 9.4 cm all along the LINAC.)

The current and pulse duration are shown in Figure 6.5.2-14; because the charge per pulse is fixed, current, and pulse duration are inversely related.

Figure 6.5.2-15 shows the pulse length as a function of displacement along the LINAC.

The total number of quadrupoles in the LINAC, derived from L(Z) and the gradient, is 878. Based on the conservative assumption that the cores surround the focussing magnets all along the LINAC, the total volume of core is 2535 m^3.

6.5.2.2 Acceleration Modules - The acceleration modules occurring in the LINAC are composed of the induction cells and the pulsed power systems that drive them. The induction cells for the single beam LINAC are relatively simple, resembling the cells of electron LINACs more than those of a multiple beam LINAC. The core volume is much smaller, surrounding only a single beamlet and focussing magnet, compared to N beamlets and magnets (see Figure 6.5.2-16).

By using the single beam LINAC, an N-fold reduction in the number of magnets is achieved, and the total core volume is reduced by $N^{1/2}$. In the LINAC the penalty paid for this simplicity is that the single core must be rapidly pulsed N times, since the beamlets are accelerated serially rather than in parallel. This increases the complexity and the cost per joule of the pulsed power.

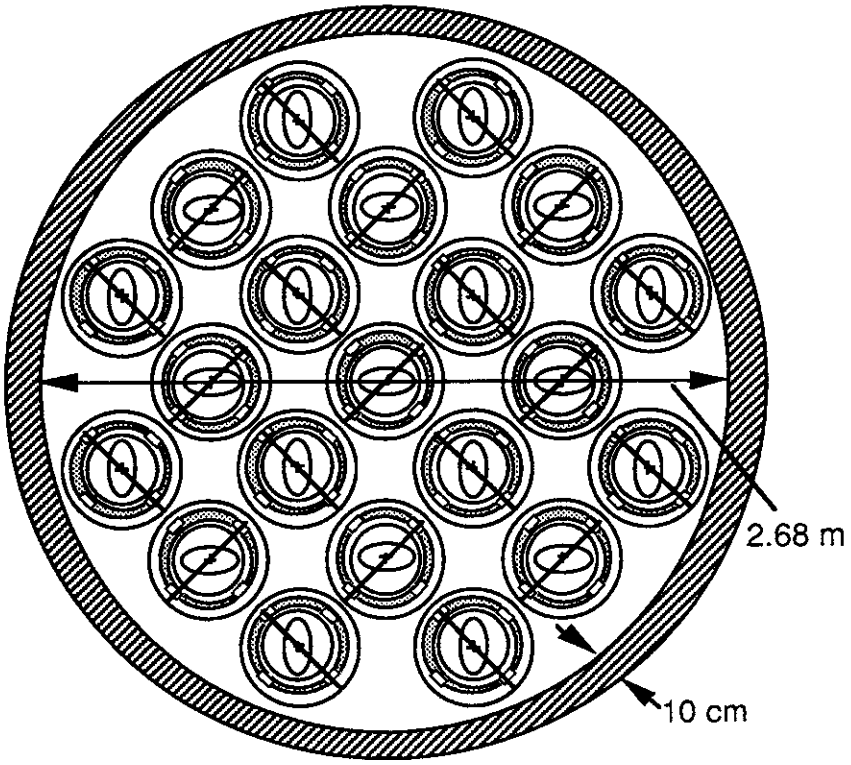

Ratio of Magnets
21:1

Ratio of Core Volumes
48:1

($\sqrt{21} = 4.6$)

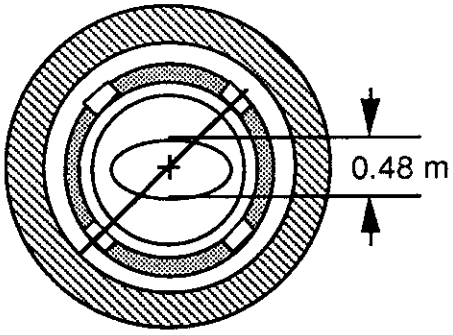

Figure 6.5.2-16 Comparison of Multiple Beam LINAC with Single Beam LINAC

Within the context of the induction cell, there are some issues that need to be addressed. First is the problem of heat generation and removal from the cores. Heat is generated in the cores due to hysteresis losses in the B-H curve and due to eddy currents in the finite resistivity Metglas material. The power loss per unit volume per target interaction is at least a factor of N higher in the case of the SBL than in the corresponding multiple beam LINAC (MBL) case, simply because the cores are being fired N times per target interaction. However, the loss curves for the Metglas core material indicate that the temperature rise caused by a burst of 18 pulses is less than 100 K, so that the ratcheting-up of the induction core temperature is still dominated by the 3.5 Hz repetition rate of the reactor. The total heat dissipated per meter per reactor

shot is nominally only $N^{1/2}$ higher for the single beam LINAC (SBL), but the cell cooling circuit must be designed to handle the higher density of energy dissipation in the smaller SBL cells.

The primary problem is the eddy current losses, which can be minimized by making the core ribbon thinner so that the skin depth is a larger fraction of the thickness and by adjusting the frequency components of the applied field. The Prometheus-H cores were assumed to be wound from 1 mil thick ribbon, rather than the 1.25 mil ribbon in common use today, and the insulator between layers was assumed to thin proportionally to maintain the same packing fraction. Metglas ribbons as thin as 0.75 mil have been produced, and plasma spraying of magnesium oxide has been suggested as a method of producing a thin insulating layer, so the modest extensions suggested are not far-fetched.

The difficulty in tailoring the waveform to minimize eddy current losses is that the time available to reset the SBL core is limited by the burst repetition rate (which is in turn set by the maximum beamlet lifetime in the downstream storage rings and the total number of beamlets). The energy dissipation for Metglas is shown in Figure 6.5.2-17 as a function of square pulse duration.

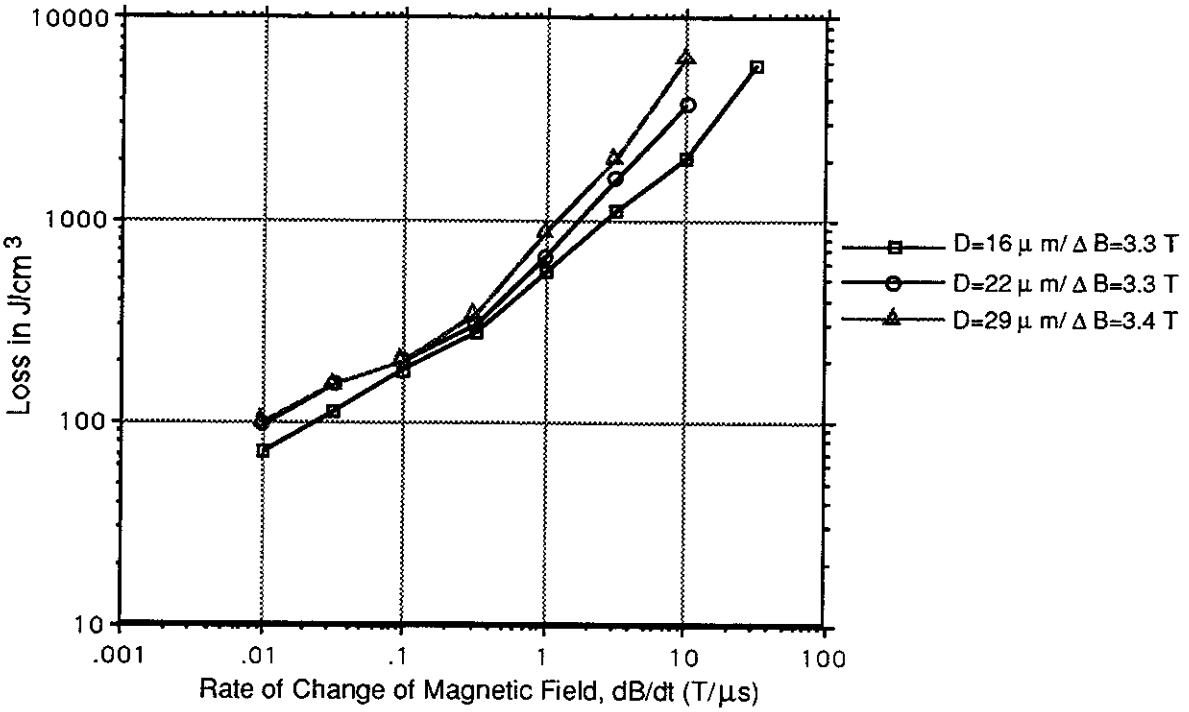

Figure 6.5.2-17 Energy Dissipation for Metglas as a Function of (Square) Pulse Duration

Real waveforms are not square and both the acceleration time and the reset time come into play, Figure 6.5.2-18 shows the heat as a function of PRF for a (1 cos θ) reset pulse.

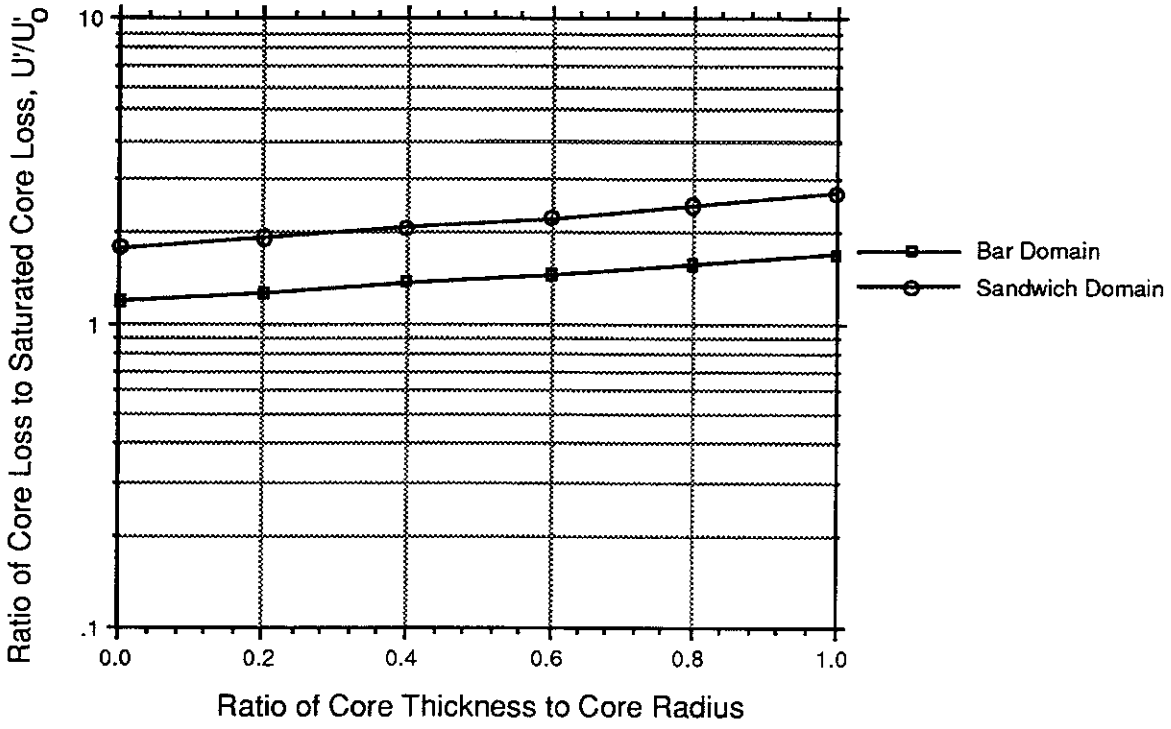

Figure 6.5.2-18 Metglas Heat as a Function of PRF

Because less time is available to reset the cores, the loss (J/cm^3) per shot in the SBL case is necessarily larger than the MBL case. Although not done for the Prometheus-H design, optimization of the waveform with respect to losses should result in lower heat generation in the core and lower installed power costs.

Another issue has to do with lifetime and reliability of the modules, since it concerns protection of the induction cell components against overvoltage in case of a misfire situation. One might think that the induction cells in a multiple beam LINAC would be less sensitive to errors such as late arrival of one of the beamlets since they are more heavily loaded by total beam current. Induction LINACs that have been built to date, including the 10 kA ATA at LLNL, have a damping resistor in parallel with the beam to limit the voltage appearing across the cell and critical components in case the beam arrival time is off (see Figure 6.5.2-19).

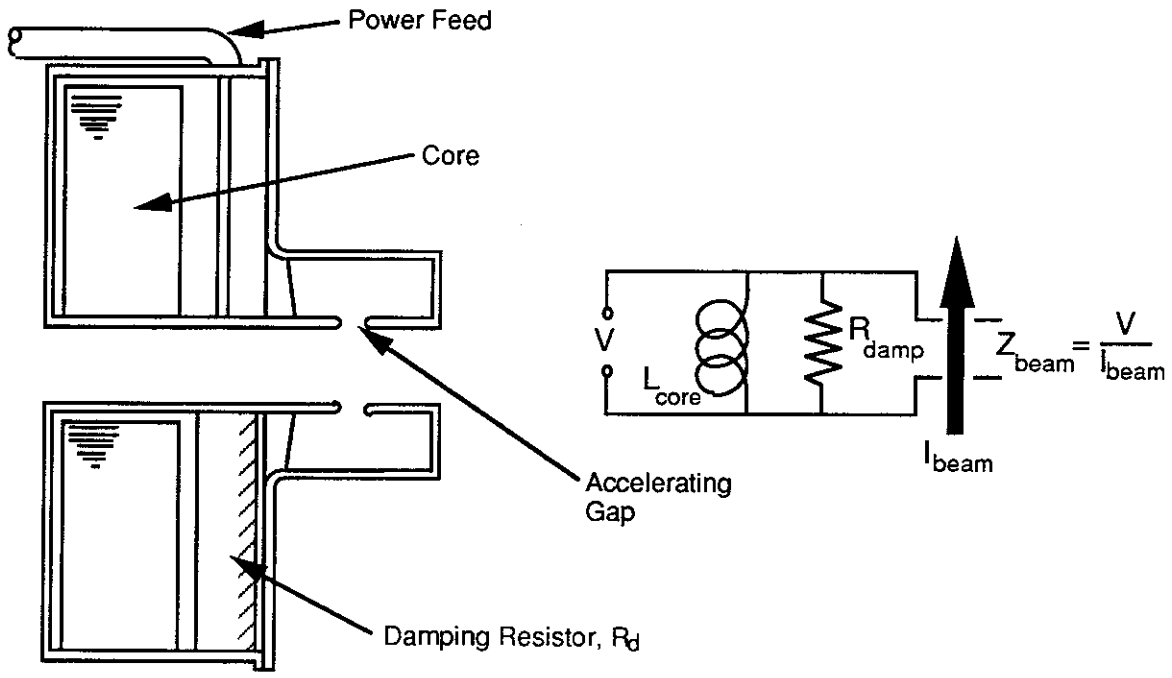

Figure 6.5.2-19 Damping Resistor Circuit on LINAC (Similar to LLNL ATA Design).

This resistor typically has an impedance equal to or even less than the beam load

$$Z_{beam}=V_{gap}/I_{beam},$$

so that the beam is only a minor perturbation to the total load. Most of the pulsed power goes into heating the damping resistor. Although this is an effective, low-cost solution to component protection in a low repetition rate experimental device, it will not be acceptable economically in a power plant simply because the pulsed power systems cost so much. Either steps must be taken to prevent firing a module in a fault condition, or the components of the induction cells in the driver will have to be designed to withstand a 2x overvoltage; similarly the pulsed power systems must be configured so they are able to withstand firing into an open circuit or be prevented from firing.

The pulsed power requirements vary along the LINAC as the pulse length, gap voltage, and current vary. The requirements for the modules are straightforward, but advancements must be made to achieve the high repetition rate burst mode operation. While thyratron switching may be used for low PRF applications such as a multiple beam LINAC, silicon-controlled rectifiers have been proposed for induction LINACs with PRF up to 6 kHz not just for pulse rate enhancement but for improved long-term reliability. The solid state devices have a near-infinite lifetime (>10^{11} shots) if they are run within their specifications, but they are very fault-intolerant: a single out-of-bounds

shot is capable of destroying the device. As a result, effort has been directed toward arranging the components so that faults in the load do not cause the pulse power components to exceed their allowed range. This evolution in configuration is illustrated in Figure 6.5.2-20.

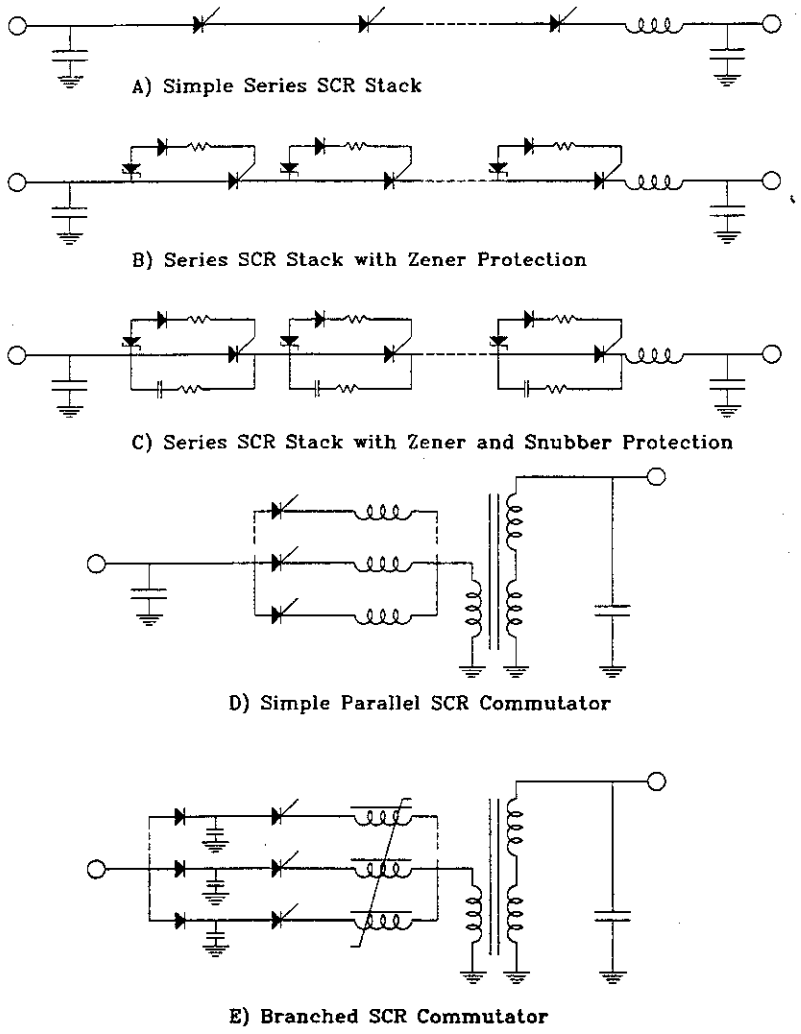

Figure 6.5.2-20 Evolved Pulsed Power Component Configuration to Limit Failures.

Work was done at LLNL and elsewhere to develop essentially cw kilojoule 6 kHz (6 MW average power) pulse compressors for use in the GBFEL and AVLIS programs, so the average power capability is there. For the factor of three higher PRF required for the SBL driver burst mode, even SCRs are probably insufficient and FETs will probably be necessary. The high pulse repetition frequency (PRF) pulsed power for the SBL will be more expensive on a per joule basis than the very low PRF used in the MBL; a value of $100/J was assumed in costing the pulsed power in the Prometheus driver.

6.5.2.3 Superconducting Magnets - The focussing quadrupoles used to contain the beam radially during acceleration will be superconducting with cos(2θ) windings to provide the quadrupolar magnetic field and covered with iron to provide flux return. These quadrupoles must fit beneath the accelerator cores (at least in the low energy end of the LINAC where the packing fraction is high) as shown in Figure 6.5.2-21.

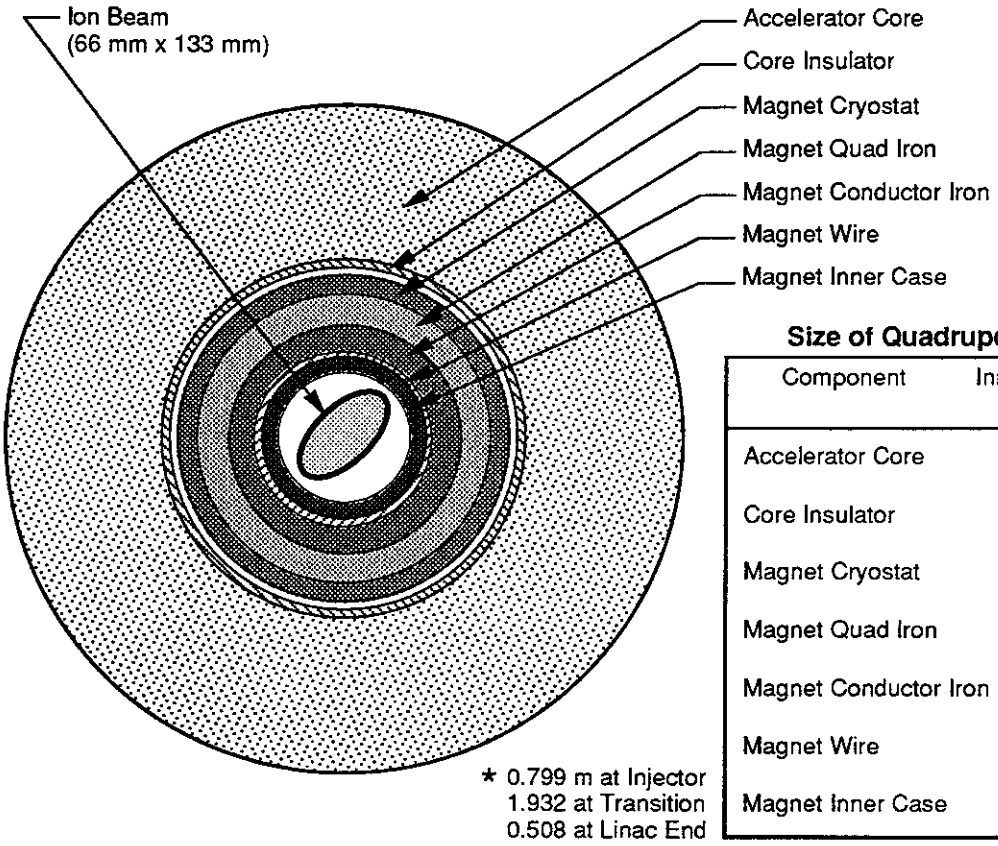

Size of Quadrupole Components

Component	Inner Radius (meters)	Outer Radius (meters)
Accelerator Core	0.43	*
Core Insulator	0.41	0.43
Magnet Cryostat	0.35	0.39
Magnet Quad Iron	0.27	0.35
Magnet Conductor Iron	0.21	0.27
Magnet Wire	0.195	0.21
Magnet Inner Case	0.155	0.195

Figure 6.5.2-21 Prometheus-H Superconducting cos (2θ) Quadrupole Core Configuration.

The cost of this type of quadrupole is relatively insensitive to the quadrupole size; however, because the induction cores fit over magnets, it was necessary to estimate the quadrupoles radial extent in order to calculate the core volume. The rule of thumb used is as follows: multiply the (circular) beam envelope by 1.65 to account for the effect the envelope ellipticity at the quadrupole location (a difference of ±30% depending on direction) and to provide a margin between the beam edge and the inner bore of the cryostat. Add 4 cm for each of two 4-300 K transitions, plus the thickness of both the superconducting windings to provide the required field and the stainless steel support structure. Finally, add the minimum thickness of laminated steel yoke to avoid saturation. Because the Prometheus-H LINAC system runs at a relatively low quadrupole gradient, the entire structure is compact.

A specific relation between the scaling law parameters was selected to provide a uniform beam diameter all along the LINAC. A consequence of this choice is that in order to meet the simultaneous constraints on axial packing fraction ($\eta < .8$) and aspect ratio ($a/L_q < 0.25$) everywhere, the ideal length of the quadrupole must vary along the LINAC (Figure 6.5.2-12).

From a spares inventory viewpoint, having unique magnets is undesirable. However, the field strength used in the quadrupoles is nowhere near the limits of the superconductors (the 2T "poletip" field in the Prometheus-H driver LINAC is less than 60% of the maximum allowed by the conservative poletip constraint). Since each magnet acts to provide a transverse momentum impulse the quadrupole length can be fixed at a standard value (or a few values) and the field varied, keeping the impulse (proportional to the product BL_q) constant.

The high allowed aspect ratio of beam radius to quadrupole length in the Prometheus-H design (1/4 versus 1/10 in other studies) is an area of potential concern. The issue is growth in the normalized emittance that is caused by magnetic lens aberrations due to the relatively short length. While there are octupole components to the lens as a result of end effects, these appear to play a significant role only if the beam is not space charge dominated. As part of the Neutral Particle Beam program, experiments have been performed at LANL on the acceleration of high current, high brightness proton beams in a large aperture drift tube LINAC. The short quadrupoles are contained within the drift tubes to keep the high current beam focussed. The beam there is observed to propagate without increasing its emittance significantly.

While it is certain that there will have to be considerable magnet design to optimize the quadrupoles, these first-cut approximations were sufficient to proceed with the Prometheus-H design.

6.5.2.4 Cooling and Cryogenics - Some observations regarding cooling not common to both types of ICF drivers are made here. The induction modules (both the cores and parts of the pulsed power) require special treatment because of the various combinations of high voltage, material corrosion, and dielectric breakdown. By their nature, the superconducting magnets require liquid helium to operate, necessitating a large distributed cryogenics system. A cryogenic system of some sort is required for the target factory as well, but the driver is spread out over a much larger area and needs a distribution system.

Induction Modules. The amount of heat generated per cubic centimeter per pulse in an induction core is small; over time the temperature does build, although it could take several minutes to become a problem. Adequate cooling in the modules and magnetic pulse compressors is a concern because the glass-like structure of Metglas makes it a poor thermal conductor and some of its properties are temperature-dependent. There must be axial gaps between Metglas pancakes for fluid, but because these gaps decrease the axial packing fraction, the core diameter must go up to compensate.

Low conductivity water (a good dielectric) is unacceptable as a cooling fluid in the cores because it will corrode the steel in Metglas. Cooling fluids such as Freon or Fluorinert (non-conductive, non-corrosive, and capable of withstanding the fields inside the module) are probably good choices as coolants. Although these fluorocarbons have other undesirable features that may force their phase-out over the next few years, suitable substitutes might be found. Their relatively low boiling points and high heat of vaporization may provide an additional advantage. For simplicity, it might be desirable to cool the induction modules passively (without forced flow). Under this scenario, the fluorocarbon is maintained at its boiling point and the heat deposited in the cores is transferred from the Metglas via heat of vaporization (forming a very small volume of bubbles). This two-phase cooling process is quite efficient. A heat exchanger then removes the heat from the fluid and deposits it in water.

The magnets in the heavy ion driver are required to be cooled to liquid helium temperature to be superconducting. The magnets are in cryostats that isolate them from room temperature. During the field ramping stage during start up, eddy currents generate a thermal load, but once operating the predominant thermal loading mechanism on the system is static heat leaks through the cryostat penetrations. Liquid helium is brought into the magnet and cold gas exits. Higher temperature gaseous helium from the boil-off is used to thermally anchor the shields at intermediate temperatures (typically 20 K and 80 K), eliminating the need for a separate liquid N2 system.

Large distributed cryogenic systems (refrigerators, cold boxes, distribution lines, etc.) of this scope have been built and operated reliably at various accelerator facilities

around the world (Fermilab, CERN, DESY, CEBAF, JAERI), and a much larger one will soon be built for the SSC. The cryogenics system is not viewed as an issue for Prometheus-H.

A properly designed cryostat will have mainly axial losses due to supports; for this reason one is driven to make the cryostat as long as possible (consistent with access requirements), minimizing the number of penetrations between 300 K and 4 K. Because the cost of a cryostat is more or less fixed, independent of size, the later portions of the Prometheus-H driver (when the beams are carried in parallel beamlines) have single cryostats containing many magnets.

6.5.2.5 Storage Rings - Each of the 18 beamlets exiting the single beam LINAC is captured and stored in one of 14 identical rings. Twelve beamlets are stored in individual rings, while the remaining six beamlets are divided into the two remaining rings. The two groups of three beamlets form the two precursor beams. The injection timing is such that the three beamlets are sequenced one after another. The precursor beams are then sent to the bunching LINACs, followed by the main beams.

Ring sizing for this study is based on a simple sublattice; focus-drift-bend-drift-defocus-drift-bend-drift. This is replicated twice and then a straight section was added to allow for injection and extraction of the beams; this was repeated four times. The storage ring is illustrated in Figure 6.5.2-22. Thirty-four 1-m long cos(2θ) superconducting quadrupoles, identical in all respects to those at the end of the LINAC, have a spacing of 4.6 m around the ring. Thirty 2-1/4 m long 6T cos(θ) superconducting dipole bending magnets have a 50% packing fraction (the magnetic rigidity of the 4 GeV Pb^{2+} ions is 65 T-m).

The storage rings share common cryostats although the beamlines are separated.

This ring is suggested only as a placeholder. No attempt was made to check the dynamic aperture even for single particles, much less a space-charge dominated beam. Nor were any aberration-correcting elements added to the lattice. Anything more serious than a cursory "reality-check" look at the storage ring is beyond the scope of this study.

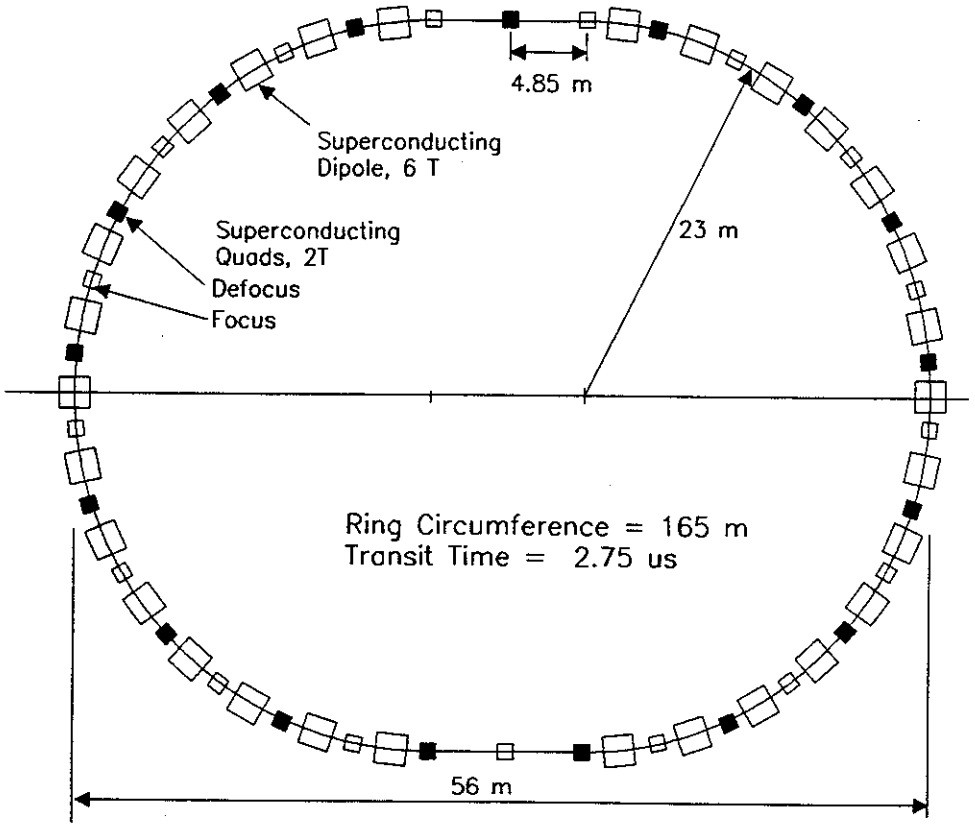

Figure 6.5.2-22 Dipole Magnet Schematic

The maximum storage time in one of these rings is 18 times the interbunch spacing. At a burst rate of 20 kHz (for example), the duration that the first bunch is stored is 50 μs x 18 = 0.9 ms. During this time, the coasting bunch makes nearly 400 revolutions, traveling almost 55 km around the ring. There is a real possibility of the beam coupling to resonances in the circular system, leading to beam degradation (emittance growth) or scrape-off. To help in this regard, the baseline burst rate of 30 kHz was chosen.

The consequences of scrape-off are more serious in a circular machine since the beam passes the same spot over and over. A slight loss of beam to the wall could evolve enough gas to further perturb the beam on a subsequent pass, eventually causing closure of the aperture through this vacuum instability; the ALADDIN storage ring at Wisconsin was plagued for years by this effect. In a linear device the gas has a chance to dissipate before the next beam pulse arrives.

At this time, all that can be truly said about the detailed dynamics within the ring is that they require a much closer look than could be given here. Storage time in the rings has been identified as a Critical Issue.

Despite the lack of a definitive ring design and ignoring the issue of whether the beam can be stored at all, some points basic to the overall concept can be addressed. Specific issues concern maintaining the longitudinal profile of the beam and coupling into and out of the rings.

First, unless prevented the beam will de-bunch as it coasts around the ring. In the central part of the beam, the axial electric field is nearly zero, but toward the ends the axial forces try to accelerate the head and decelerate the tail. To remain at the space charge limited current, the beam will occasionally have to be "kicked" at its ends to counteract the cumulative effects of the axial space charge force. The beam head must be retarded and the tail accelerated. This can be done with very few cores per ring. Because the kicker voltage is low and the affected portion of the beam is small, there is little stress on the induction modules. The core volume is negligible. This adjustment does not have to be made often, perhaps less than once per revolution. The process is illustrated in Figure 6.5.2-23.

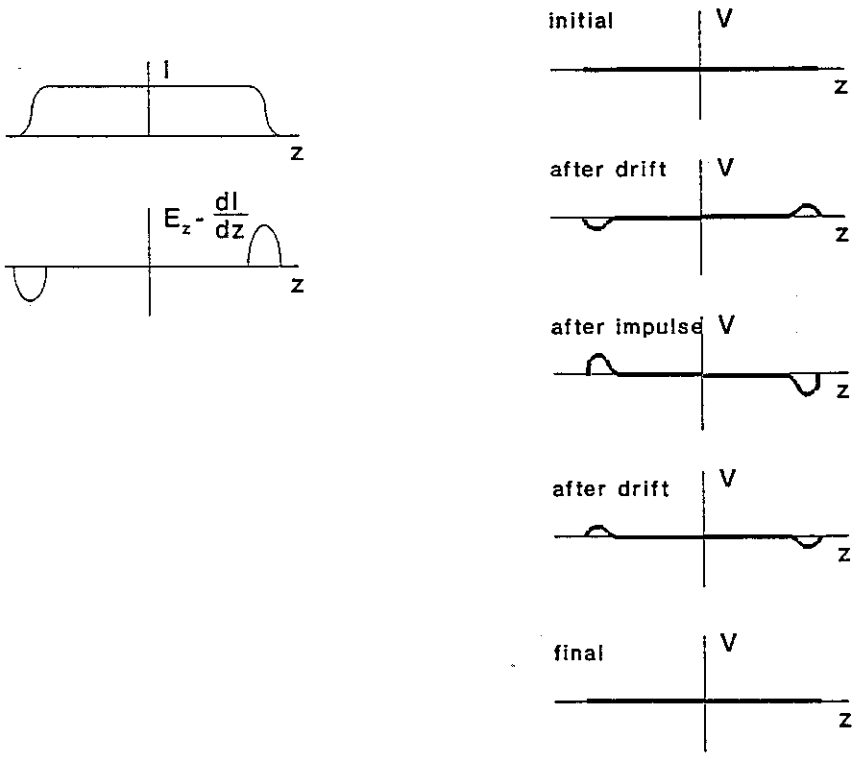

Figure 6.5.2-23. Effects of Drift and Overcorrection on Longitudinal Phase Space

In fact, this energy adjustment must be made to the heads and tails of the beam throughout the driver. In the main LINAC and the buncher, this voltage can be provided by the induction modules already present. Separate correction modules are required only in the storage rings because the beams coast at constant energy and constant length and there is no need for the main modules. However, the short length of the correction pulse may make it desirable to use separate function kicker modules in the rest of the driver as well.

It has been predicted that a high-current, heavy ion beam will couple to the impedance of gap structures in the LINAC, forming a longitudinal analog of head-to-tail cumulative beam-breakup in electron LINACs. Small-scale variations in beam energy along the pulse (with wavelength long enough that the space charge forces do not reduce them quickly) would be transformed to current variations by the drift, and it is easy to imagine an instability that would grow to some finite level before it saturates. The outcome would be to increase the longitudinal emittance of the pulse and make it more difficult to compress. The problem would presumably be worse in a storage ring, where the total path length is much longer and the possibility of a regenerative instability also exists. Mechanisms for control of the instability include lowering the impedance of the gap and using feedback in the cells to damp the modes.
A method must be provided to assure that the beams in all rings are synchronized so that they are at the correct position when they are ejected from the ring. Small differences in beam energy or ring circumference could otherwise accumulate, leading to a condition in which the beams are separated in time by as much as 2.3 μs due to one beam passing the extraction point before the others. RF storage rings use the principle of phase stability to maintain the particle in the right position; a slight periodic ramp over the pulse would accomplish the same thing in the long-pulse induction systems, accelerating late pulses and retarding early ones. The parameters would have to be controlled carefully.

As an aside, the previous two points serve to illustrate the need for a feedback control system, not just for the storage rings but for the entire driver. Diagnostics to sense the beam and mechanisms to control it must be integrated into the driver. An expert system will probably be required to make Prometheus-H feasible. Interest in model-based systems is increasing at accelerator facilities as the machines become larger and more complex. It is probable that operations knowledge and experience accumulated in the R&D efforts in support of the heavy ion driver will be used in educating an expert system.

Finally, there is the problem of injection and ejection. In the high current ring, the quadrupoles are necessarily close together, leaving very little room to couple into and out of the ring. Consider Figure 6.5.2-24, a detail of the ring described above. There are only 3.7 m of space between quadrupoles with 0.43-m radius. Bringing the beams out between adjacent quadrupoles does not appear to be practical. A pulsed

dipole magnet producing 10 T would have to be 2.7 meters long in order to allow the beam to clear the quadrupoles. The magnetic field energy alone is on the order of 40 MJ per magnet per pulse; even assuming that the energy is not dissipated, the field would have to be established within 2 μs.

A more practical solution might be to use an approach suggested by the LLNL team working on recirculating induction LINACs. Several large-bore quadrupoles would be inserted into the lattice, and the beam would be slowly extracted through them via multiple kicks. This concept is illustrated schematically in Figure 6.5.2-25.

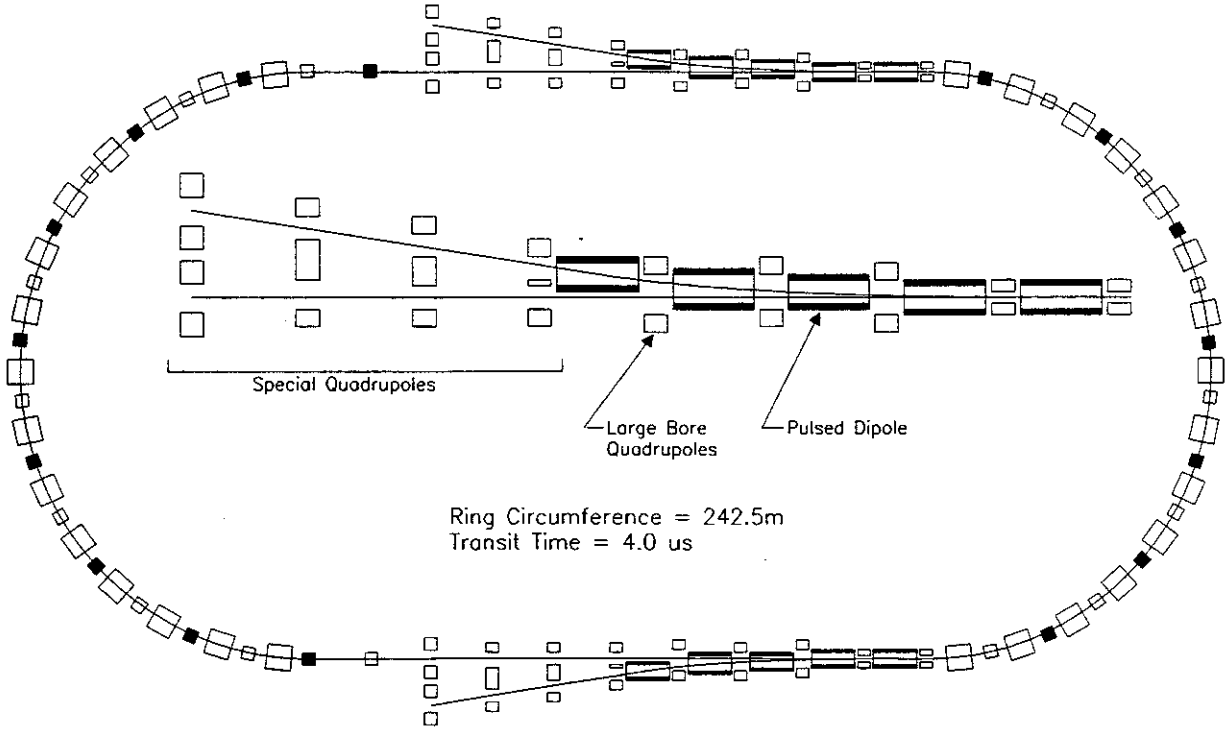

Figure 6.5.2-24. Portion of Ring with Gap

The issue of matching into ring lattice from the LINAC and into the final focus from the rings must still be addressed.

A number of issues may have to be resolved before the community accepts the idea of a single beam LINAC and high current storage rings as a viable alternative to a multiple beam LINAC (MBL). The MBL approach requires the development of very large induction modules and very complex multiple-channel superconducting quadrupoles. The single beam LINAC (SBL) approach shifts the technology onus from the LINAC, which now becomes within reach of demonstrated technology, to the storage rings. However, the potential cost savings achievable with the single beam system may be the factor that determines whether an HIF power plant is eventually built.

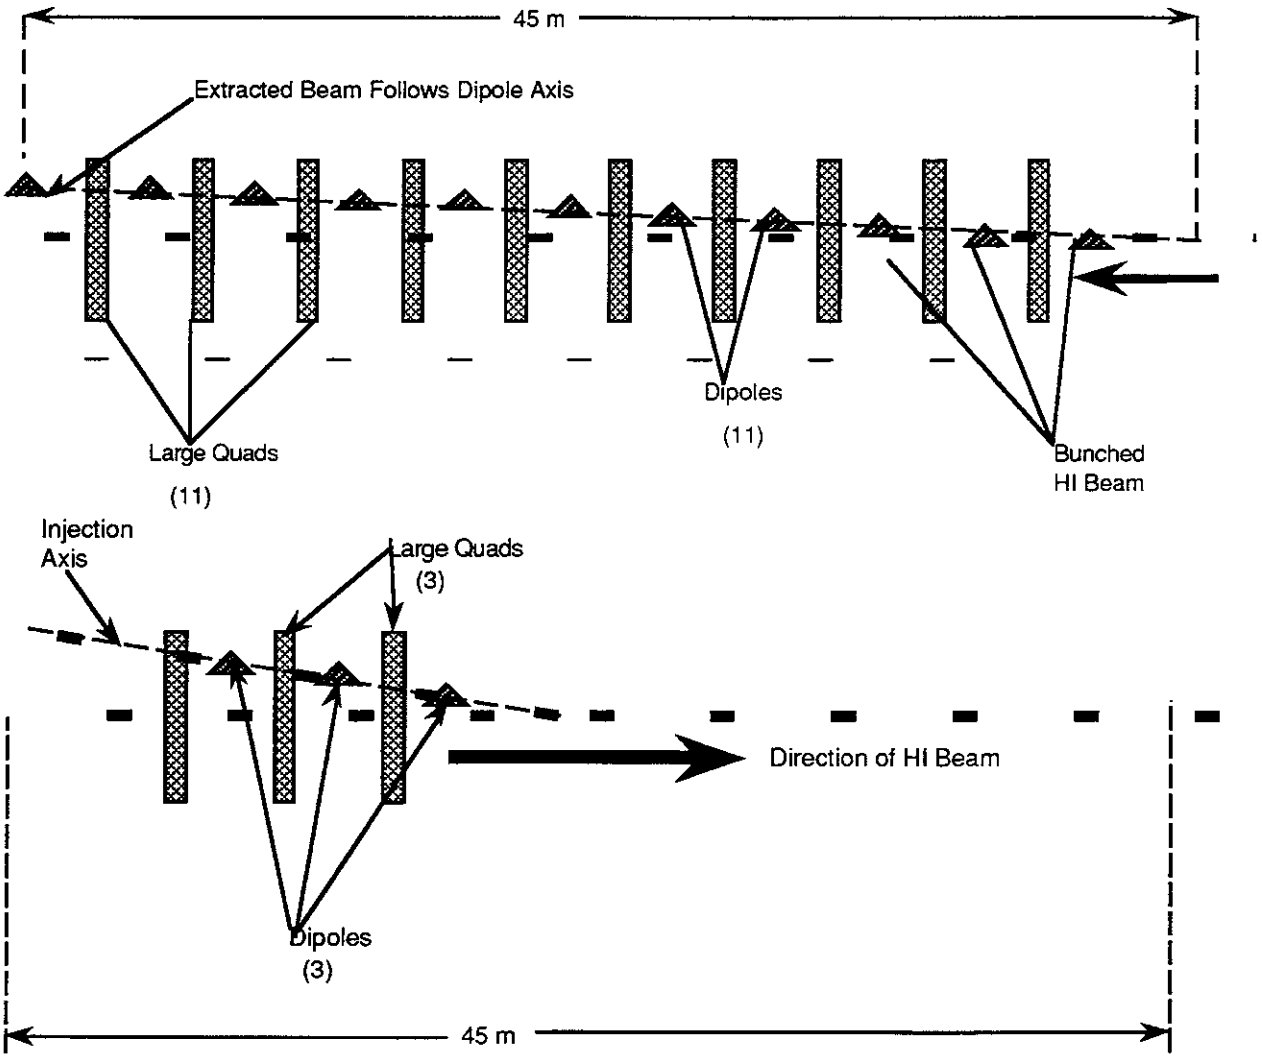

Figure 6.5.2-25. Schematic of Multiple Quads and Beam Ejection Techniques

6.5.2.6 Final Transport - The final transport of the beams to the target consists of several steps. The majority of the beamlets are bunched, transported to the reactor area in two groups, tightly focused, and finally combined in a small spot just outside of the target chamber. The remaining beamlets are combined before focusing and form the two precursor beams with a longer duration and a lower total energy content. The beam is transported to the target in a self-pinched mode through a preformed channel.

Bunching - When the storage rings eject the 14 pulses, the beamlets are sent to one of two induction-LINAC based bunchers. The first buncher compresses the twelve main beams, the second compresses the two precursor beams. The two bunchers are in close proximity, sharing the same tunnel. Both bunchers are configured with independent beamlines, with superconducting quadrupole focusing magnets sharing common cryostats and the induction module cores surrounding the beamlines between cryostats. The configuration is illustrated in Figure 6.5.2-26.

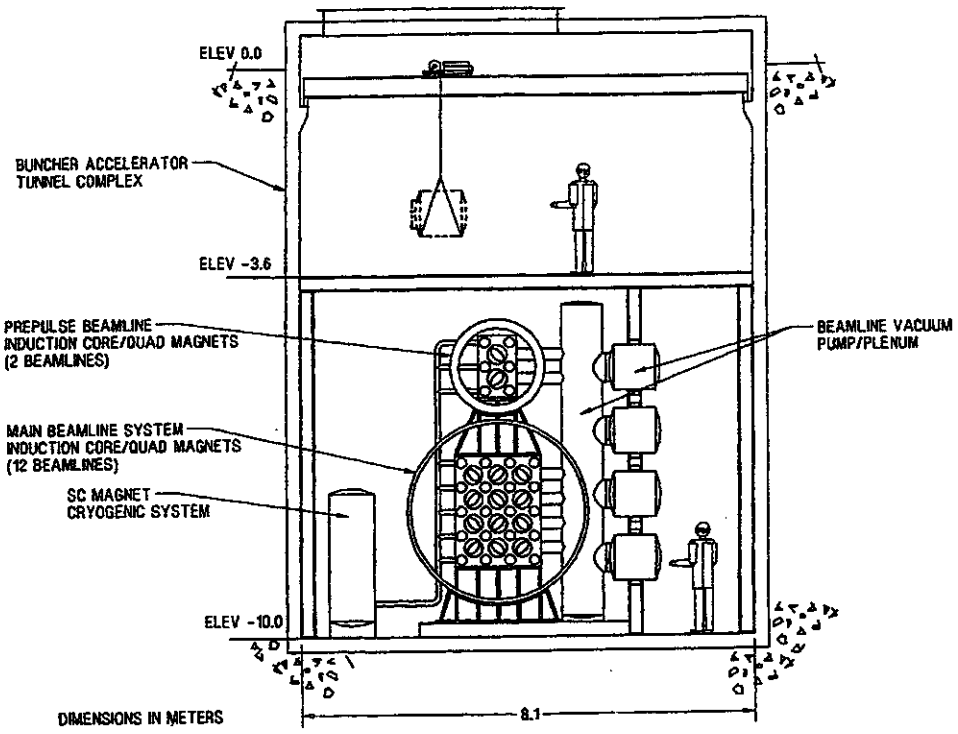

Figure 6.5.2-26. Precursor and Main Pulse Bunchers Equipped with Common Cryostats

The quadrupole lattice parameters at the entrance to each buncher are the same as at the SBL exit.

Each buncher LINAC adds a position-correlated energy across its beam pulse. The buncher cells are ramped in time, adding zero to the front of the pulse and producing the maximum 1 MV/m gain at the end of the pulse. The main (precursor) buncher extends 85 m (180 m), so there is a 4.2% (9%) energy difference from the front to the back of the pulse. The bunchers do add to the total beam energy, but ignore that for the purpose of this discussion and consider the beam to be coasting. In practice the energy of the LINAC would be dropped 2.1% to compensate for the net energy gain in the buncher. The precursor beams would still be 2.1% high in energy and the precursor buncher would be readjusted accordingly.

As the bunches drift, the backs of the pulses begin to catch up with the fronts, increasing the current. This bunch compression process starts within the buncher and continues up to the target. The beam current is shown as a function of distance from the front of the buncher to the target for both the main and precursor beams in Figure 6.5.2-27.

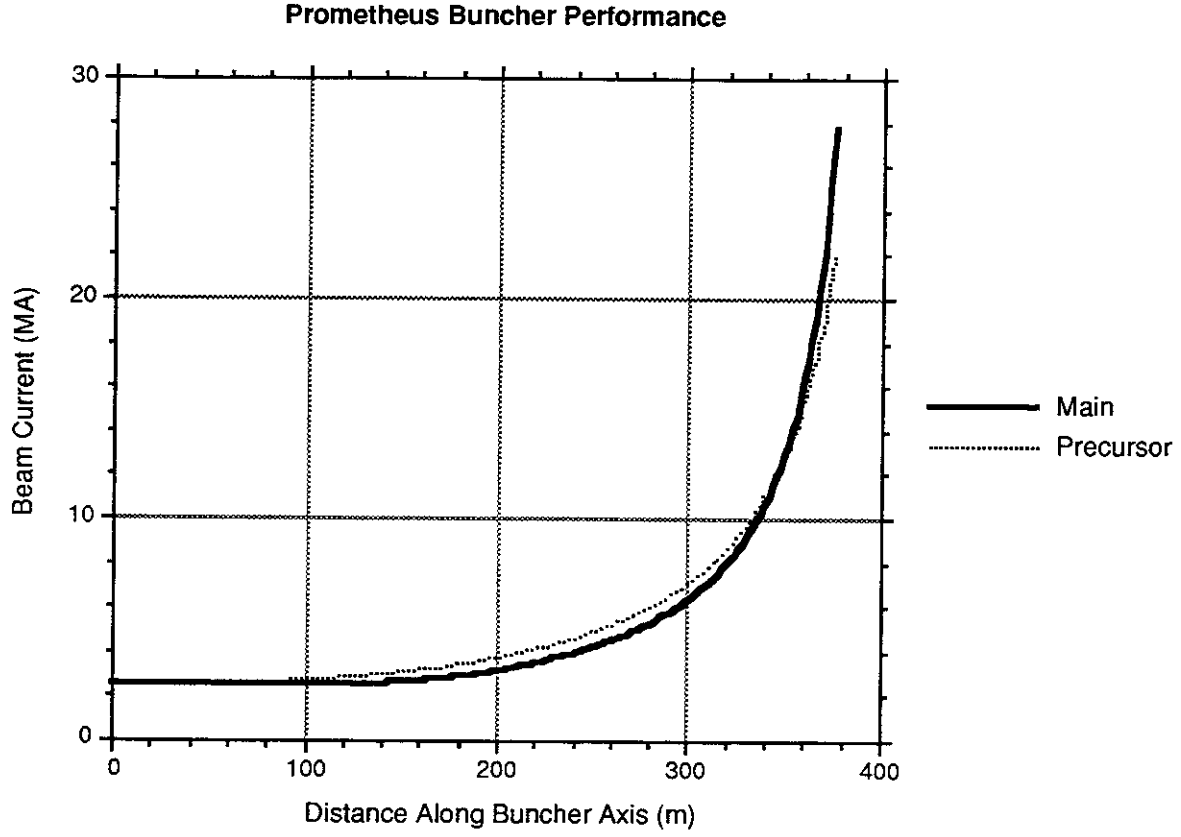

Figure 6.5.2-27. Plots of Beam Current for Precursor and Main Beams vs. Distance Along Accelerator Axis

Note that the length of the buncher tunnel is dominated by the precursor, not the main pulse; the main pulse actually coasts for 80 m in the tunnel before the bunching begins.

Like the rest of the system, the final transport is adjusted to keep the beam at its space charge limit, even as the current is increasing. The driver scaling parameter used in the constant voltage portion of the accelerator was $\gamma = 0.8$. The focusing lattice parameters are a function of the local current, and the current profile along the final transport was used to produce curves of lattice parameters versus distance shown in Figures 6.5.2-28 through 6.5.2-34.

These curves terminate about 30 m short of the target since, at that point, the focusing channel approach is abandoned and the beams enter the final focus.

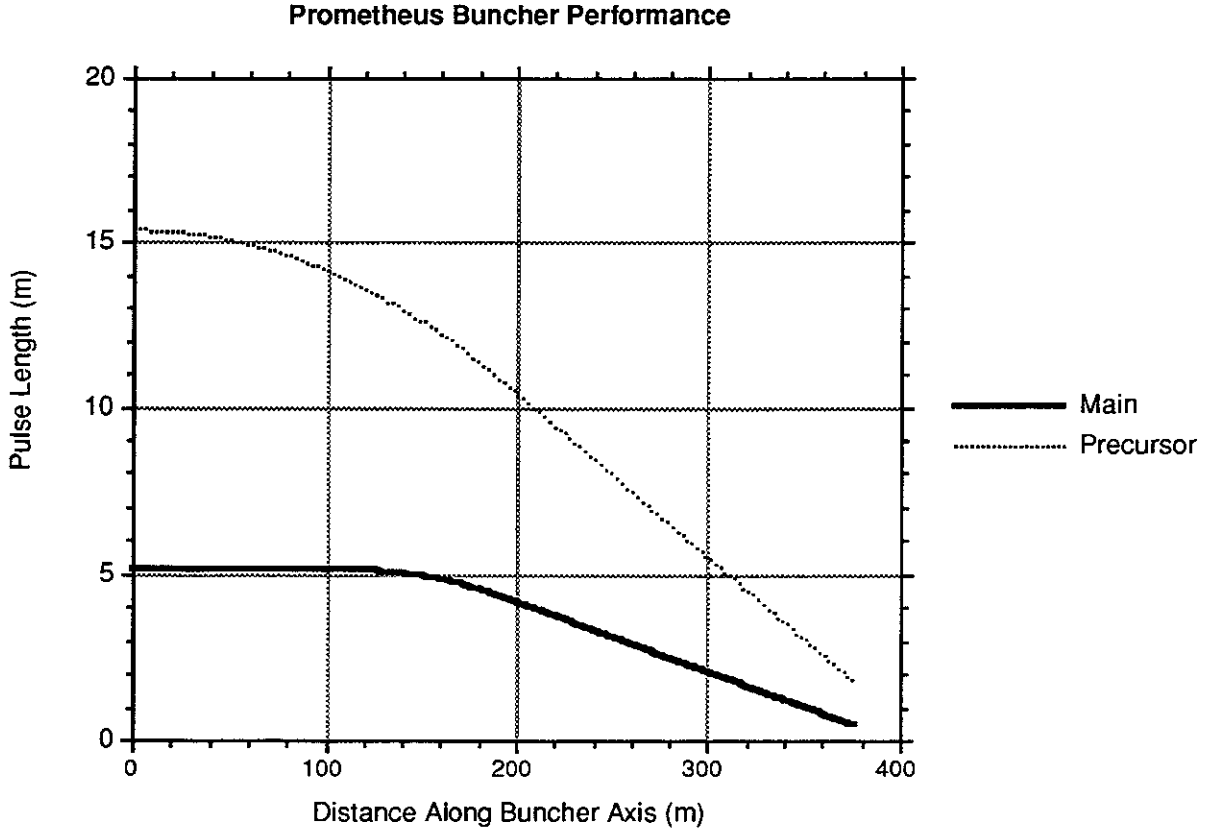

Figure 6.5.2-28 Pulse Lengths of Main and Precursor Pulses vs. Distance Along the Buncher Axis.

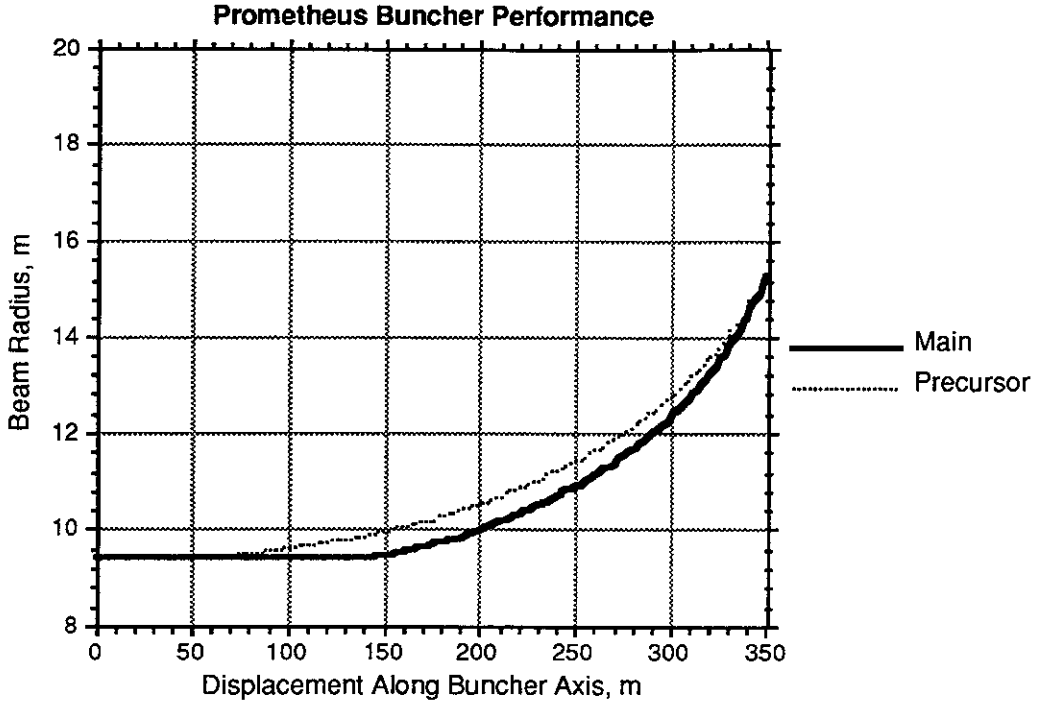

Figure 6.5.2-29 Beam Radii of Main and Precursor Beams vs. Distance Along the Buncher Axis

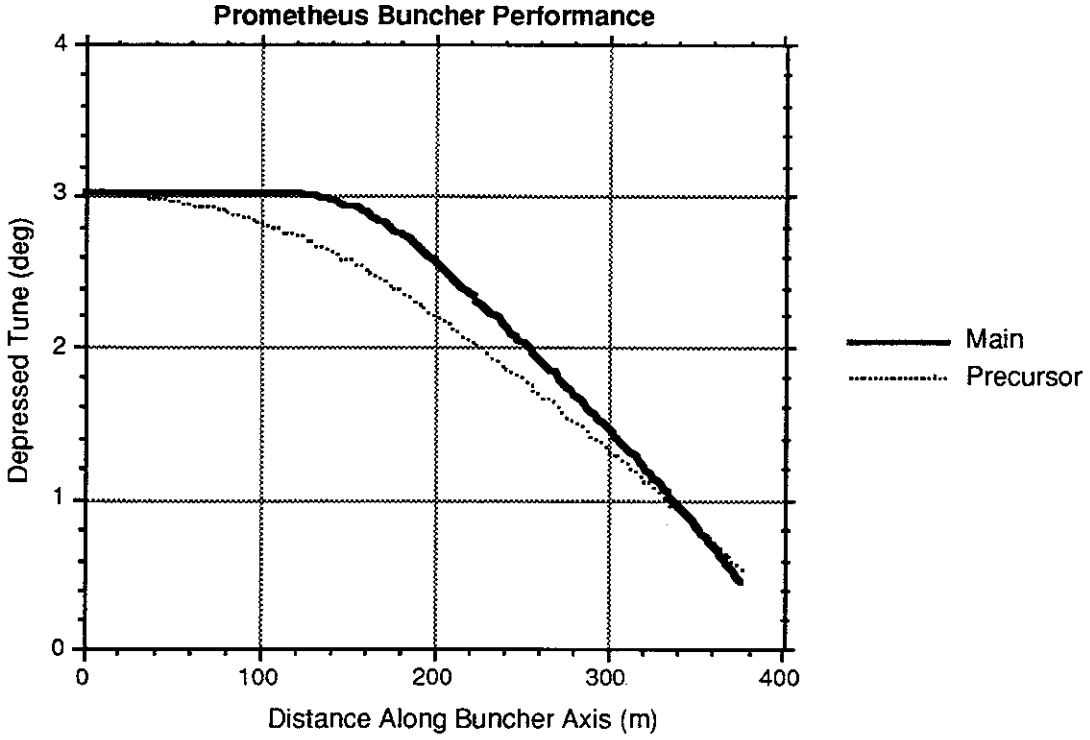

Figure 6.5.2-30 Depressed Tune as a Function of Distance Along the Buncher Axis

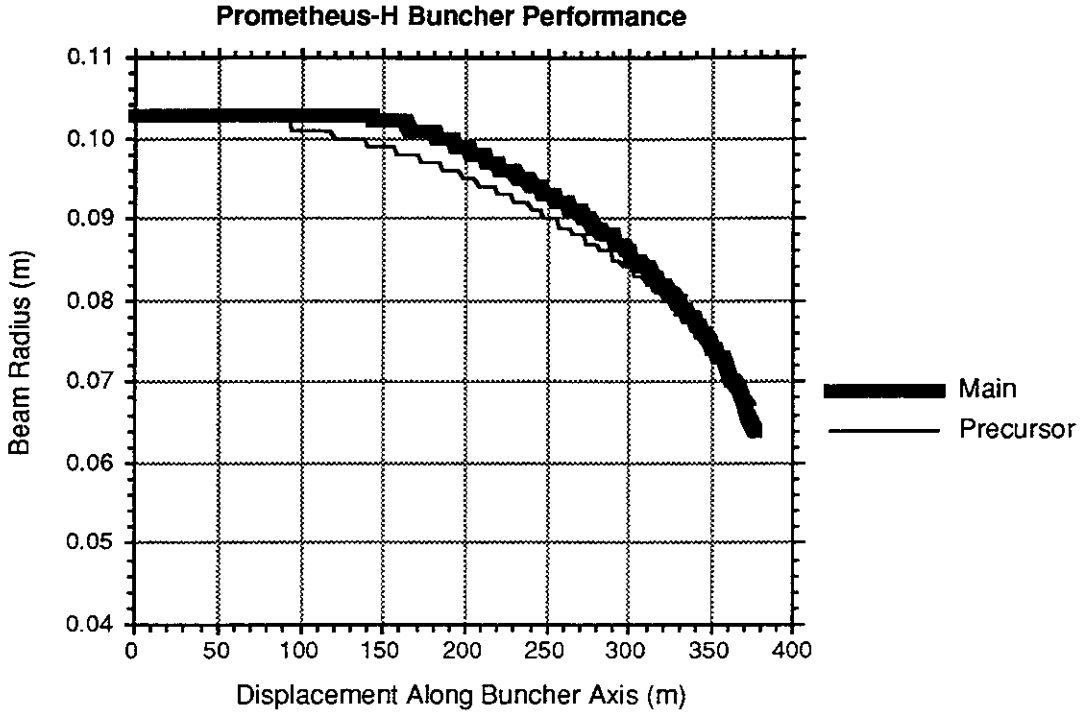

Figure 6.5.2-31 Ratio of Beam Diameter Over Length as a Function of Distance Along the Buncher Axis.

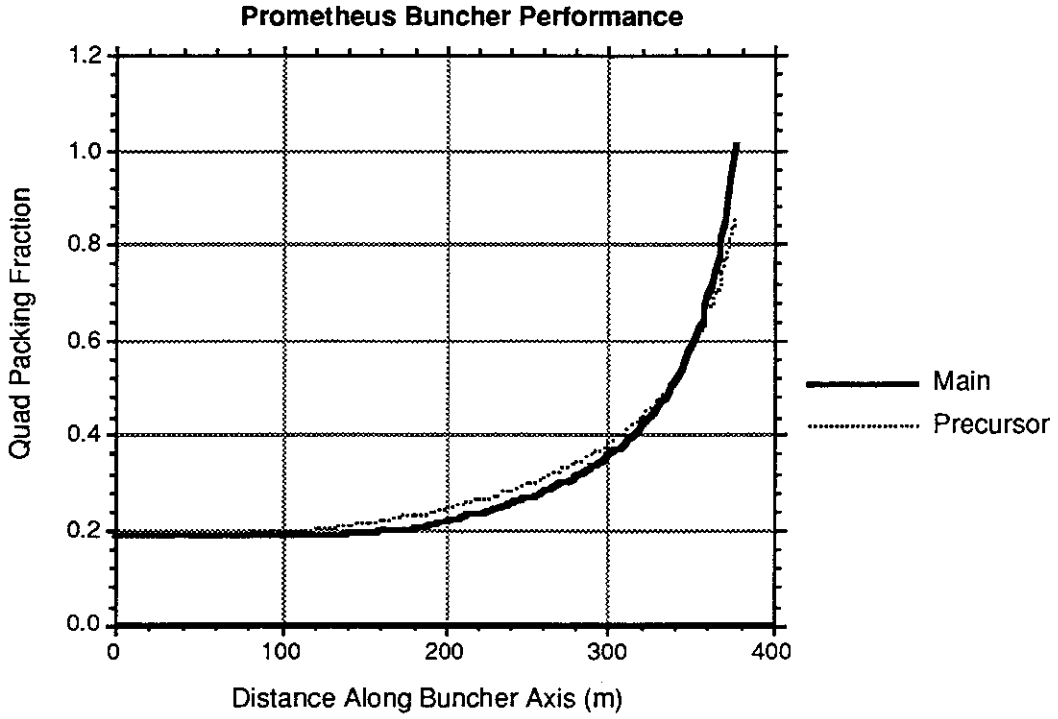

Figure 6.5.2-32 Quad Packing Fractions of Main and Precursor Beams vs. Distance Along the Buncher Axis

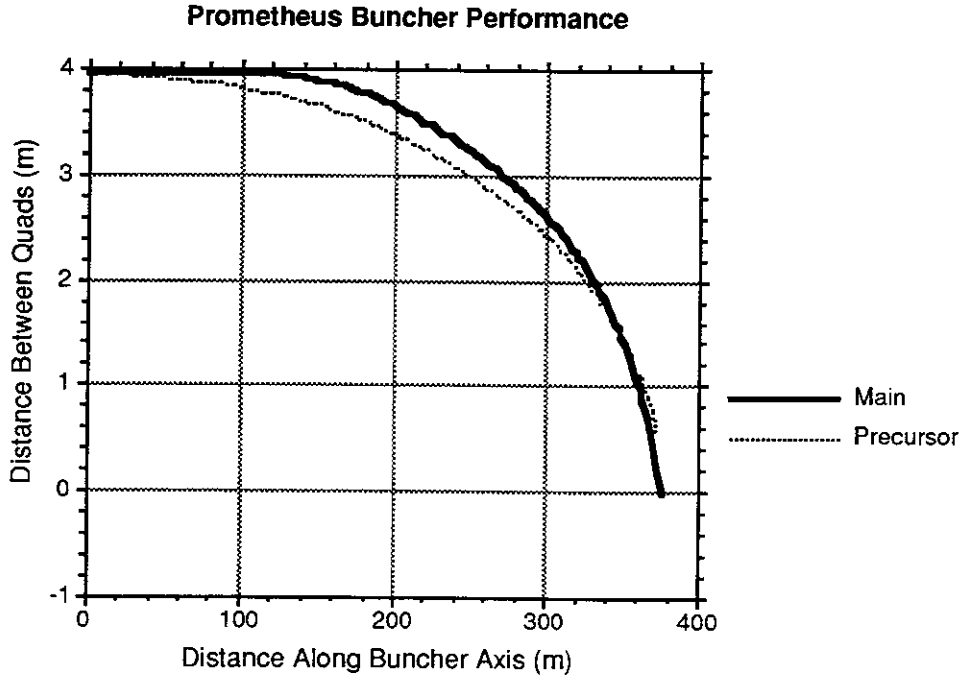

Figure 6.5.2-33 Distance Between Quads for Main and Precursor Beams vs. Distance Along the Buncher Axis.

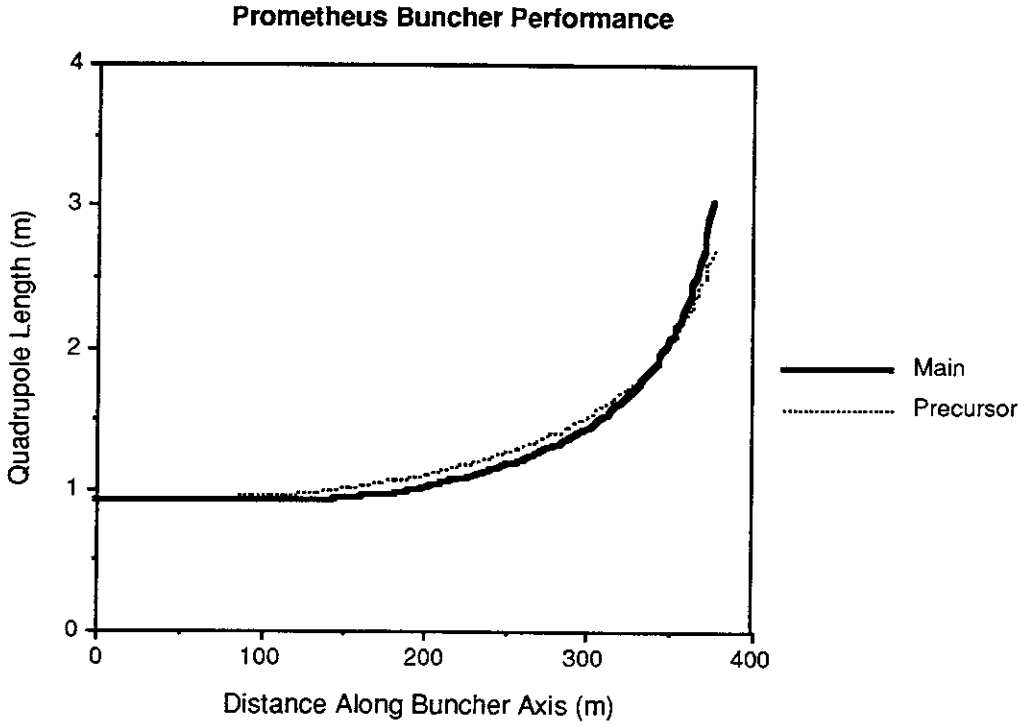

Figure 6.5.2-34 Quad Lengths for Main and Precursor Beams vs. Distance Along the Buncher Axis.

The same constraints on packing fraction and aspect ratio as used in the rest of the LINAC were used in the final transport calculations. Toward the end of the drift, the beam radius, and hence the model's "ideal" quadrupole radius, increases significantly. As was the case in the LINAC, the quadrupole strength will be adjusted so that a standard quadrupole is used.

At the end of the bunchers, half the beams (six main and one precursor) go to one side of the reaction chamber and half to the other. For most of the distance, the six main beams are bundled to share a common cryostat, but as they approach the reactor chamber building, they first diverge and then regroup as they go into the final focus. On each side, the precursor beam comes in on axis and the six main beams are arrayed symmetrically about it.

As was stated at the outset of this section, the transport schedule is very aggressive. Near the final focus section, the depressed tune has dropped to one degree, one-eighth of its value at the injector. While no instabilities have been identified, the system is on the very edge of stability. The betatron period is longer than the distance from the storage ring to the target. Extreme care must be taken not to significantly perturb the beam since there is virtually no restoring force.

Final Focus - The assumptions made about the lattice concerning packing fraction, etc., are violated before the target is reached. The beamlets are released from their FODO arrays before the limits are reached. The beams expand radially under the influence of their space charge as they drift toward the final focus magnets. The beams are reformatted with quadrupoles before they enter the final lens system. A quadrupole triplet is the simplest lens system that focuses in both planes and is astigmatic to lowest order. The magnets in this quadrupole triplet are low field superconducting cos(2θ) coils, similar in principle to the quadrupoles everywhere else in the driver, but larger and containing shielding against the neutrons from the reactor. The magnets are illustrated in Figure 6.5.2-35.

The triplet is adjusted to focus a monoenergetic, low-charge beam to a 6-mm diameter spot a distance approximately 10 m from beyond the last magnet (Figure 6.5.2-36).

The magnets are sized so that the outer diameter is as small as possible and the magnet closest to the reactor defines the solid angle subtended by the triplet as seen from the focus. All seven lenses are tightly packed (Figure 6.5.2-37).

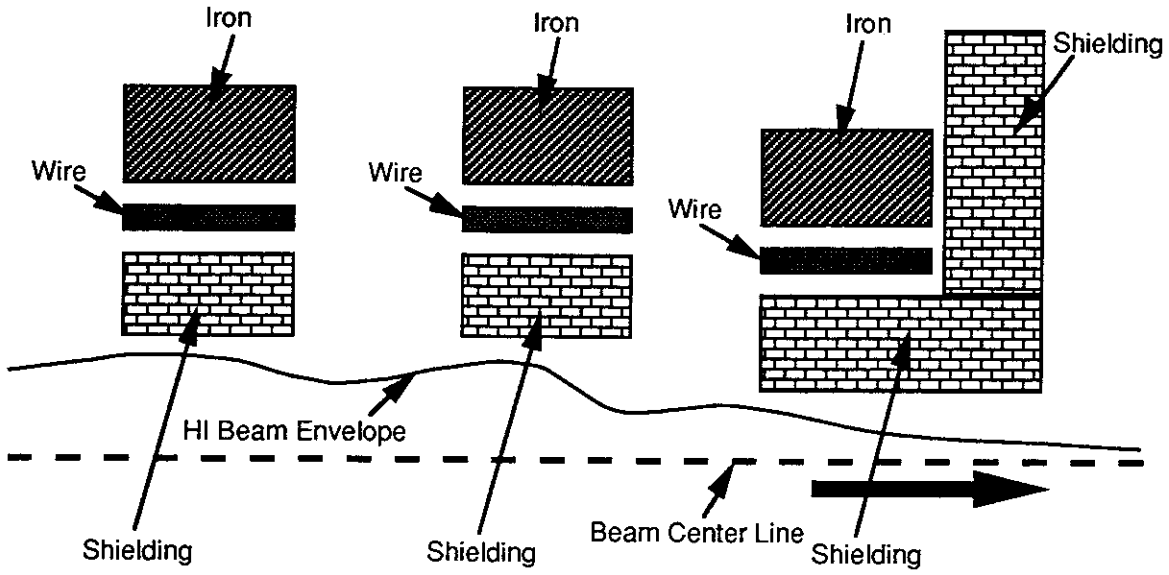

Figure 6.5.2-35 Quadrupole Triplet Magnet Cross-Section.

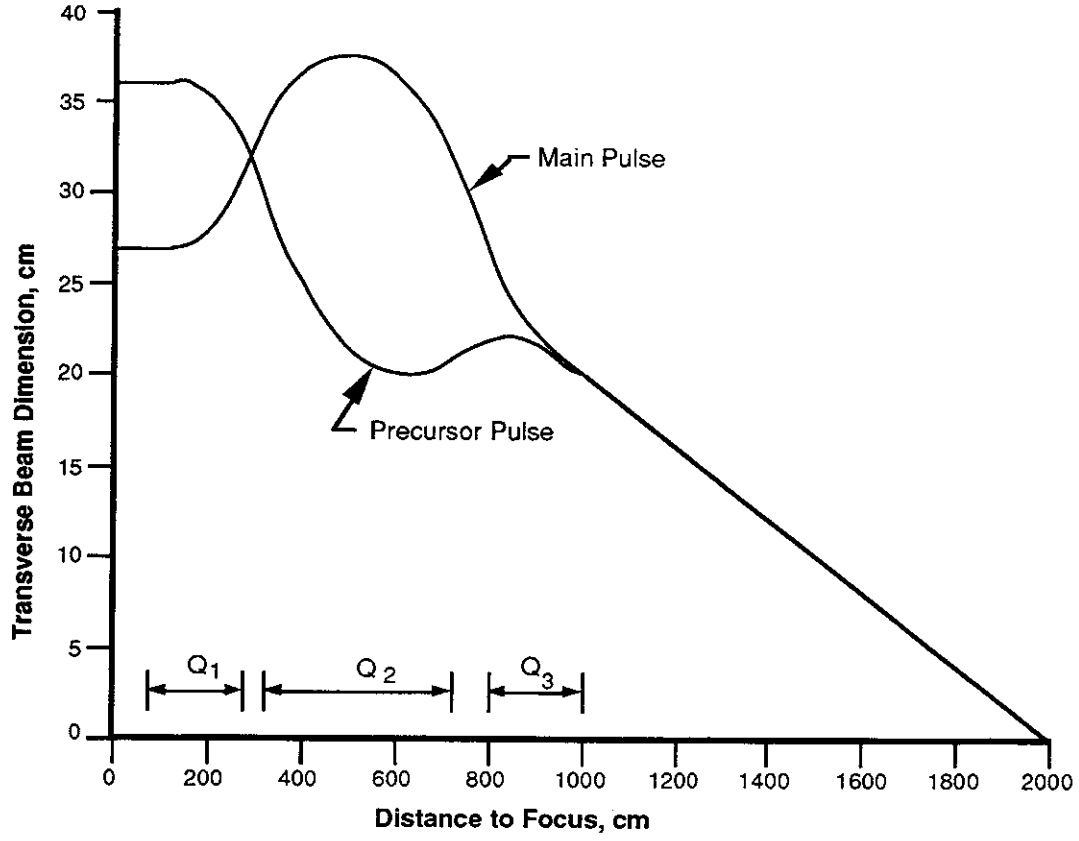

Figure 6.5.2-36. Convergences of Precurser and Main HI Beamlets in Final Focusing

The allowable f/# of the lens (focal length divided by effective aperture) is limited by geometric aberrations on the low side, and by beam emittance on the high side. However, the incoming beam is neither monoenergetic nor neutral and lens chromaticity and space charge of the beamlet become problems before these limits are approached. The approach taken in the Prometheus-H driver final focus design was to have the final focus dominated by emittance of the beam and work to reduce the effects of space charge and momentum spread on lens performance.

Chromatic Aberration - Even without additional effects related to space charge and high current transport, a simple magnetic lens is not achromatic. The focal length of a quadrupole is proportional to the momentum, and a spread in beam momentum translates directly to a spread in image distances from the lens. In a plane at a fixed distance from the lens, the apparent beam spot size changes with momentum.

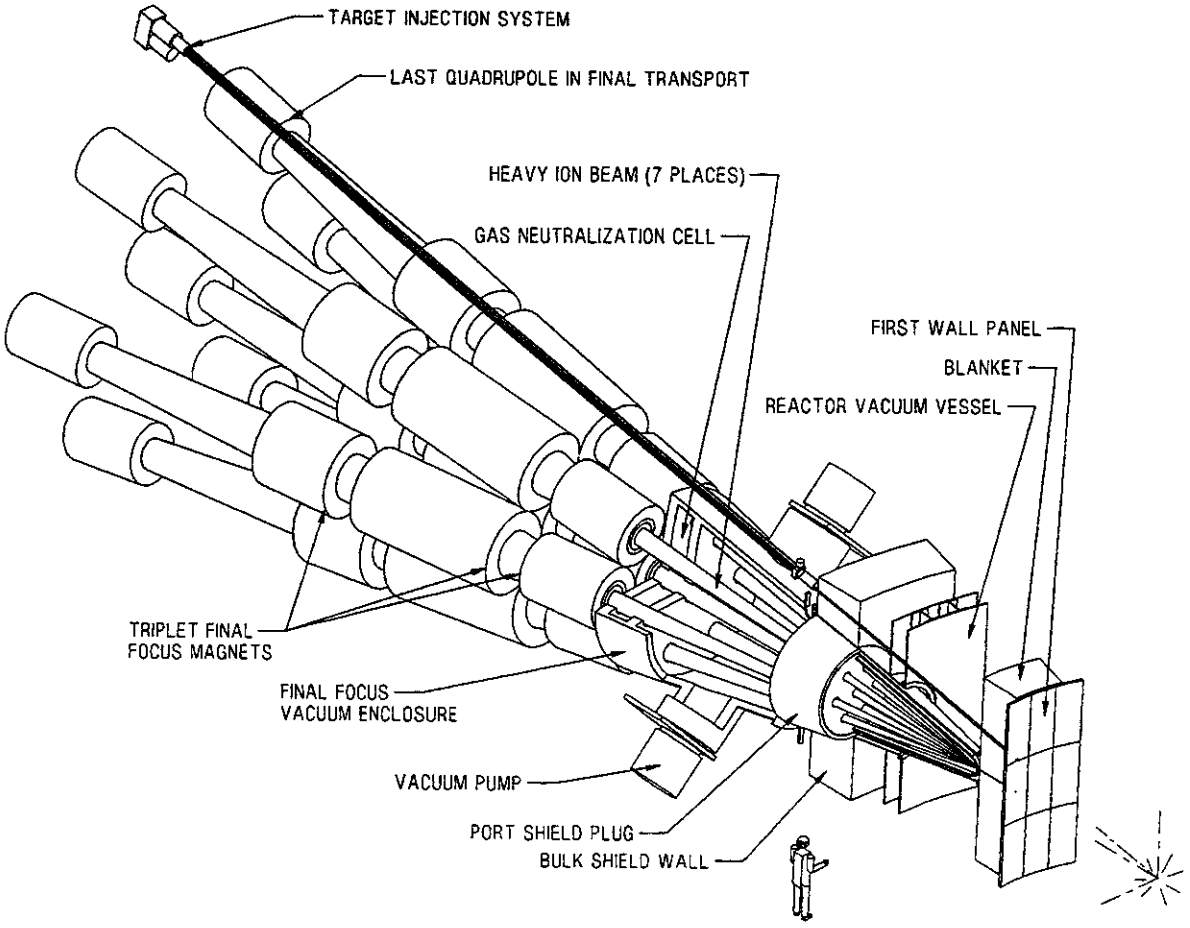

Figure 6.5.2-37. Final Focus Physical Arrangement of Final Focus Triplet Coils, Vacuum Plenum, and Shielding

The inclusion of a pair of dipoles was suggested in the HIBALL-II study to add compensating position-energy dispersion to make the final focus achromatic to first order, as shown in Figure 6.5.2-38. The spot diameter at the target position is then nearly constant. However, the residual dispersion causes the focal spot of a 1% spread beam to sweep across the target in one dimension (as shown in Figure 6.5.2-38). A net improvement is realized by adding the dipoles, but the gain is not all that one would hope.

Longitudinal Tailoring. There have been simulations appearing in the literature showing that, by precisely tailoring the axial spatial and momentum profile upstream, the beam will compress longitudinally in such a way as to completely eliminate energy spread at the lens position (as shown in Figure 6.5.2-39). Residual chromaticity causes beam with 1% energy spread to sweep across target, lowering time-averaged power on target.

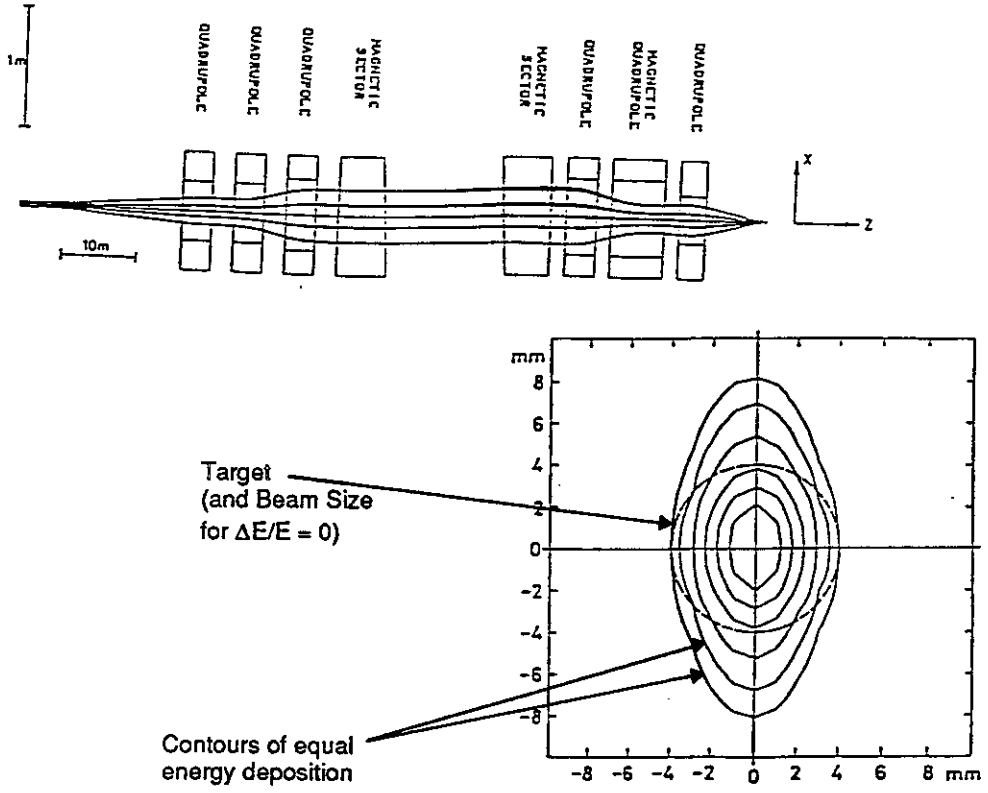

Figure 6.5.2-38. Chromatic Aberration Effects with Magnetic Dipoles.[1]

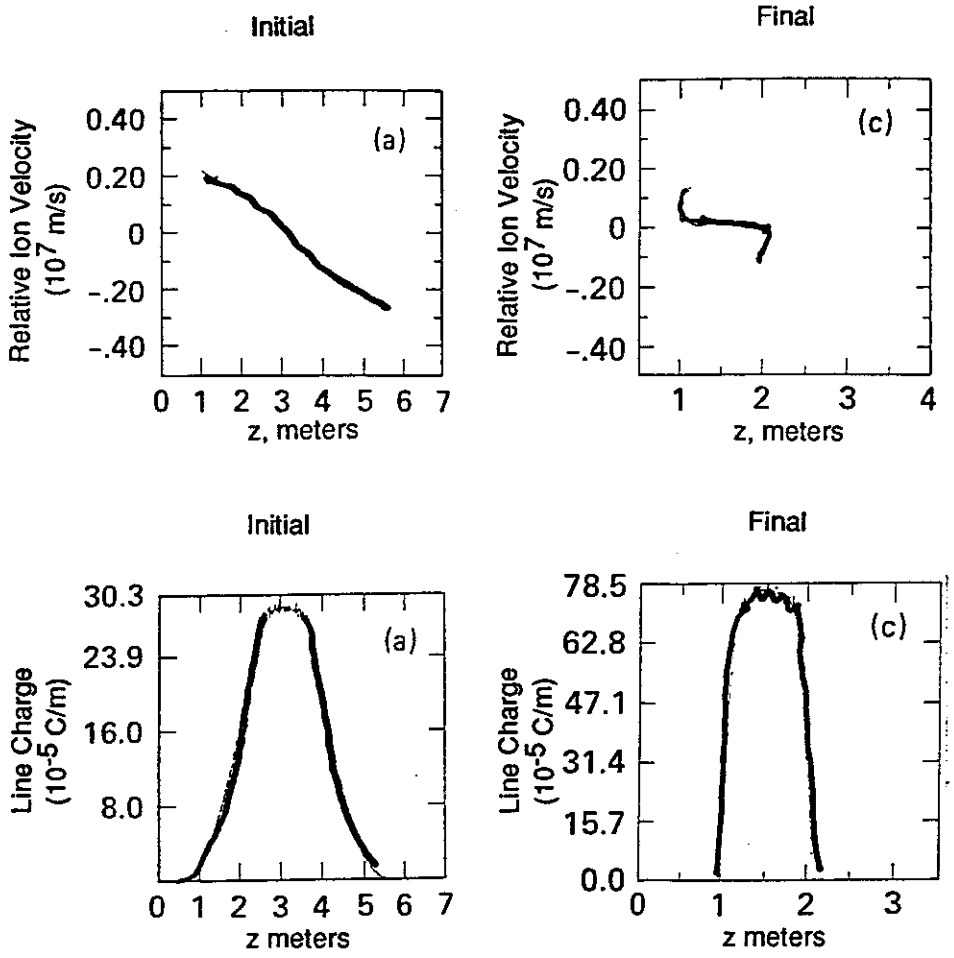

Figure 6.5.2-39 Longitudinal HI Beam Compression of a Monoenergetic HI Beam[2]

The beam reaches its peak current at the final lens; the current then drops on its way to the target as the beam re-expands under space charge forces. In practice, the degree to which the axial momentum spread can be reduced at the lens is uncertain.

Transverse Tailoring. A second approach, developed for use in the Neutral Particle Beam program, is to keep the image point of the lens fixed by altering the apparent object point of the lens to track the changing focal strength. This method relies on correlating the instantaneous transverse properties of the beam with the instantaneous momentum. The place to do this is at the exit of the buncher, where there is a nearly perfect correlation between position in the pulse and the momentum. Relatively low field, slowly varying quadrupole lenses are used to shift the final focus lens object position to compensate for the momentum-correlated focal length shifts of the final focus lens. This is illustrated in Figure 6.5.2-40.

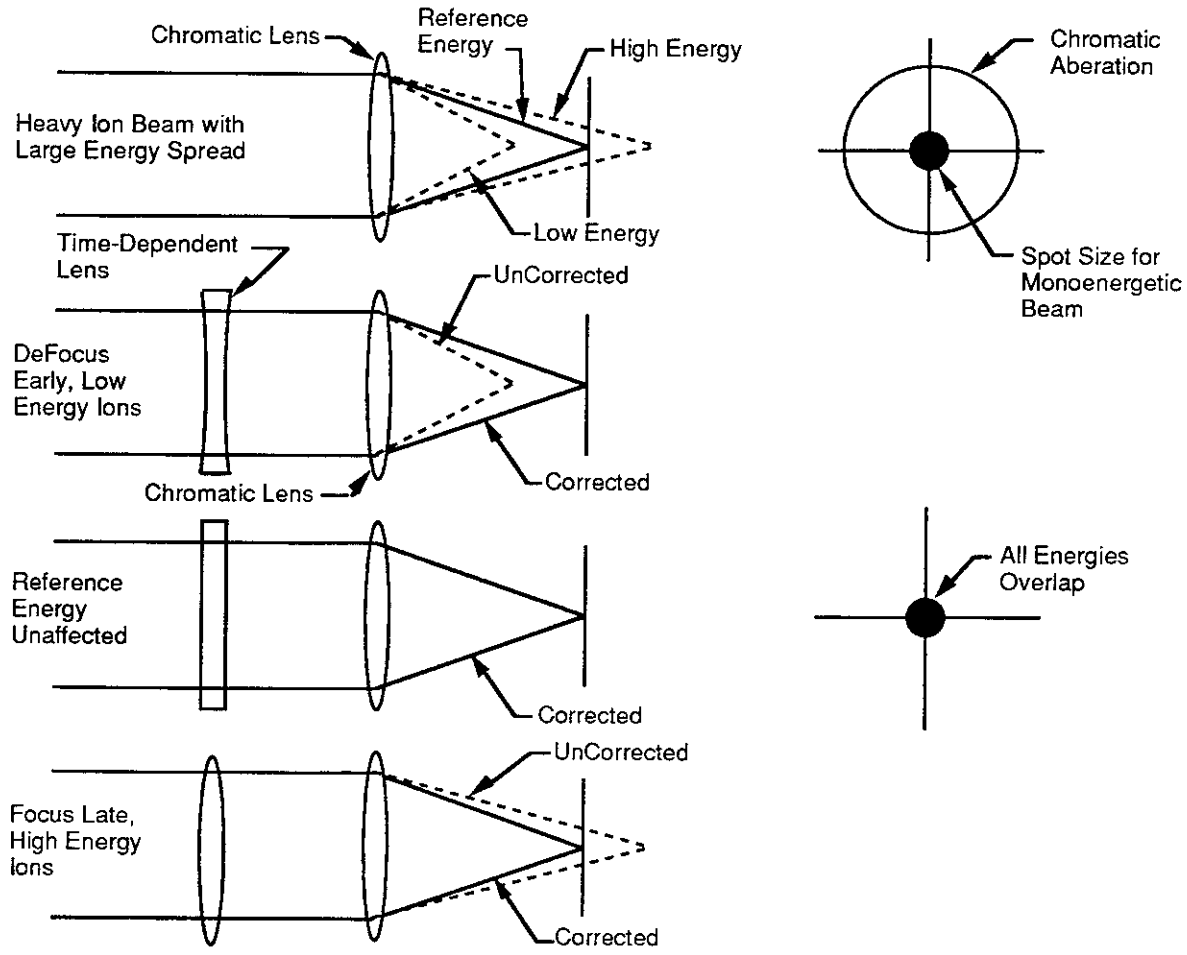

Figure 6.5.2-40 Optical Focusing Analogies for Triplet Quadrupole Focusing of HI Beams.

Simulations show that this approach works as well as the longitudinal tailoring. An additional advantage is that the beam is still contracting when it passes the final focus lens; the current peaks at the target rather than the lens. For this reason it was chosen for the Prometheus-H baseline.

Space Charge Neutralization - Beam space charge is an annoyance within the triplet lens, since it is an additional radial de-focusing force. Note that, because the beam is still contracting at the final triplet, the space charge is a function of position within the lens system, although still much lower than the ultimate current at the target. The individual magnet strengths can be tweaked to compensate, although the degree of adjustment will be larger for the main beam than the precursor.

Past the final lens, space charge is a major problem because it dominates the final spot size. By time-reversing the radial free expansion of a beam (in a manner similar to that discussed under the transverse tailoring section above), we see that the

minimum convergence angle required to reach a 6-mm diameter spot 10 meters away is 0.6 mR. If the beam is to reach its required diameter, the effects of beam space charge must be eliminated soon after it leaves the final lens.

Channel Formation - A number of methods to propagate a high current, high voltage, heavy ion beam to the target have been suggested,[1] as summarized in Figure 6.5.2-41.

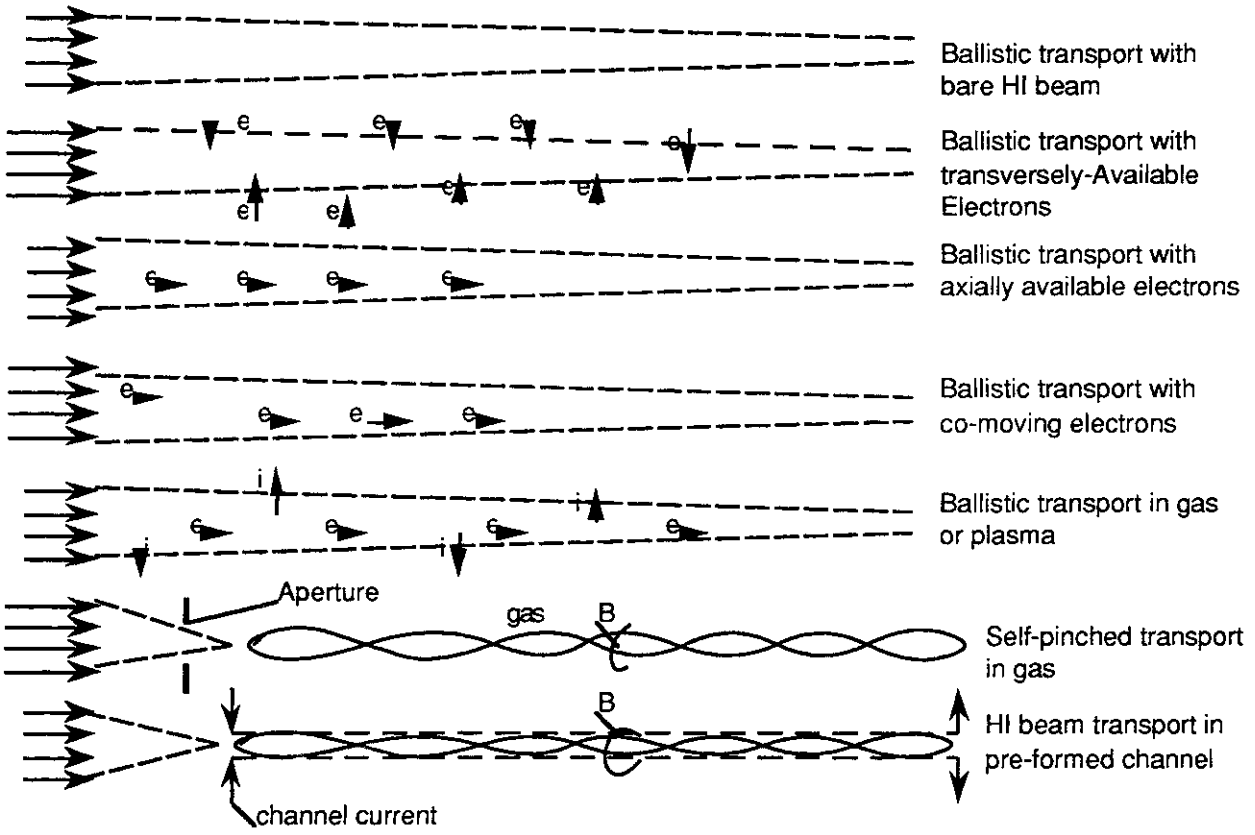

Figure 6.5.2-41 Techniques for Propagating HI Beams to Targets.[1]

With the exception of propagation through hard vacuum, they all rely on plasmas and are all speculative to one degree or another. Detailed mechanisms of propagation, as well as susceptibility to a host of instabilities, depends on the plasma's response to the beam. Most of the guidance as to allowed parameter ranges comes either from analytic models or from computer simulations coupled to atomic physics theories or measurements. Depending on the details of the assumptions, different results are obtained and there is nothing to anchor the codes. Some data comes from experiments with low energy, high current light ion diodes. There appears to be no experimental data on anything approaching a relevant regime for a heavy ion driver, and there won't be any data until a device is built that can approach the required current and energy.

However, given that propagation through hard vacuum does not lead to attractive reactor designs, one is forced to make a choice among the various plasma-propagation methods. The beam can be neutralized by adding bare electrons directly to the beam (e.g. via field emission or thermionic emission) or by adding a plasma whose density is much higher that of the beam. In the former case, it is difficult to imagine a practical scenario (one compatible with a reactor) that injects electrons uniformly into the beam. In the latter case, the beam is not necessarily completely neutralized, but its effective charge is reduced by the ratio of beam density to plasma density since the electrons act to shield the beam ions as well. Due to the inertia of the plasma ions, they can be considered stationary for the duration of the beam pulse. If the electrons remain cold (the DeBye length much smaller than the beam size), the neutralization can be quite effective.

Passive (collisional) ionization of a gas by the beam and other autoneutralizing processes tend to generate co-moving electrons that drift along at about 1±.3 times the beam velocity. Simulations indicate that the electrons heat adiabatically as the beam compresses radially by nearly two orders of magnitude. These heated electrons lose their capacity to shield the beam ions, which then see their own space charge and defocus again. In order to create enough electrons to neutralize, the gas density is fairly high. Some stripping of the beam can occur, raising the ion charge state. As long as the beam ions are neutralized by cold electrons, the electric field within the beam is very nearly zero and the ion charge state does not matter. As the shielding deteriorates, the ions that have higher charge states are defocused more strongly. This is the option adopted in the Prometheus design.

If a lower density gas is preionized, the plasma can provide sufficient electrons to neutralize the beam, without further ionizing the beam to any great extent. With a source of fresh cold electrons available, electrons that begin to be heated by radial compression can escape from the beam, leaving the shielding effective.

In another option, the drift space following the lens is filled with low density gas which is then preionized before the beam arrives. This is a transient process in which a small amount of gas is introduced via a fast valve and then ionized with a pulsed RF burst. From a practical standpoint, pulsing the plasma is necessary to vacuum-isolate this final transport region near the reactor from the UHV region upstream. On the time scale of beam propagation through the region, the gas is stationary; the beam is gone before the gas reaches the magnets. On the time scale of the reactor shots, there is minimal gas loading and the gas is gone before the next pulse 200 ms later. This may also be a viable option.

With each beamlet transported ballistically through the plasma, all eventually converge on a 6-mm diameter spot just outside the blanket, about 5 m from the pellet.

The precursor beam arrives first, traveling on axis. All main beams overlap spatially at this spot with an 8.5 degree angle between individual beams and the axis.

Transport to Target - At the point where the beams converge, there is a 2-cm diameter cylindrical hole passing through the blanket. The hydrogen partial pressure in the blanket is approximately 100 mtorr (see Table 6.8.2-9 for details) and the hole in the blanket is the major gas load for the final focus region. To lower this load, a thin, high velocity jet of lead vapor is directed across the aperture. This jet forms a gas barrier to the target chamber, entraining the escaping hydrogen in the flow like a diffusion pump. Since the lead is condensible, it does not constitute a vacuum problem even if it escapes the jet. A detail is shown in Figure 6.5.2-42.

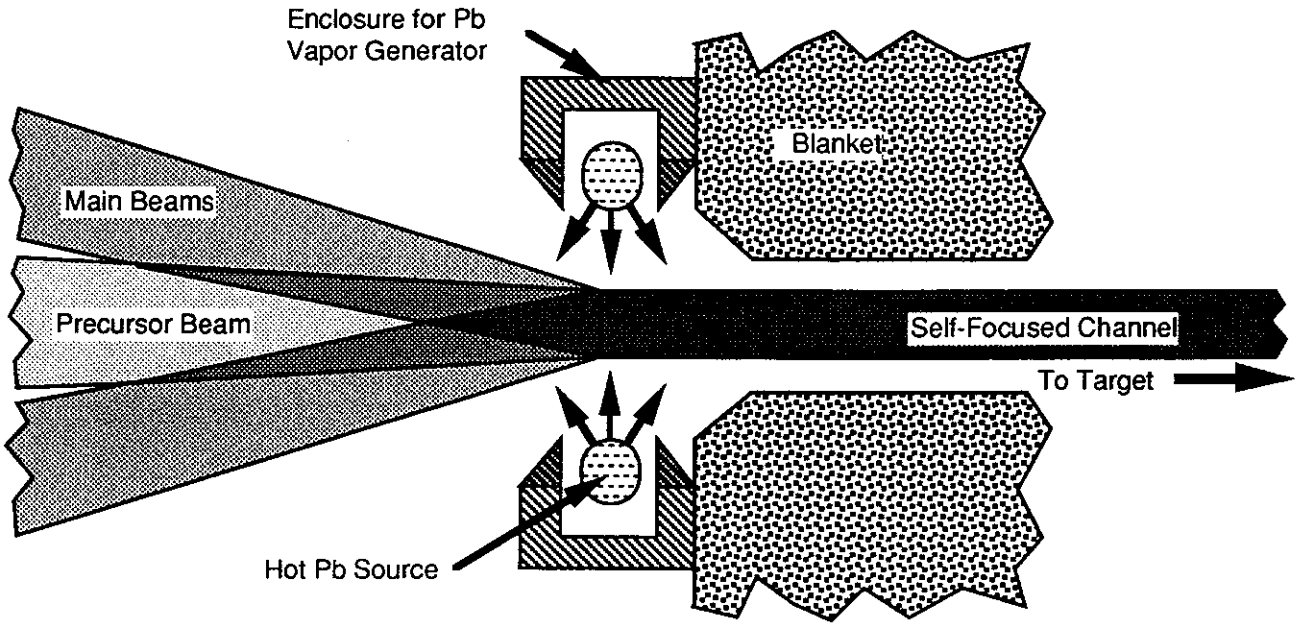

Figure 6.5.2-42. Schematic of Lead Jet System for Stripping Pb Heavy Ions

More importantly, the heavy ion beam is completely stripped in this lead vapor jet. The charge state increases from 2+ to 82+ and vastly multiplies the ion beam current. If it were not for the current multiplication, the beam would diverge into separate beamlets on the other side of the jet and collide with the blanket. Instead the current is high enough that the beam pinches and propagates 5 m to the target without expanding.

The initial direction of pinched beam propagation is presumably defined by its net momentum. The precursor is aimed along the axis and, assuming axisymmetry, it should go where it is meant to. The six main beams are merged and the net momentum, if the beams are perfectly balanced, is directed along the axis as well (see Figure 6.5.2-43).

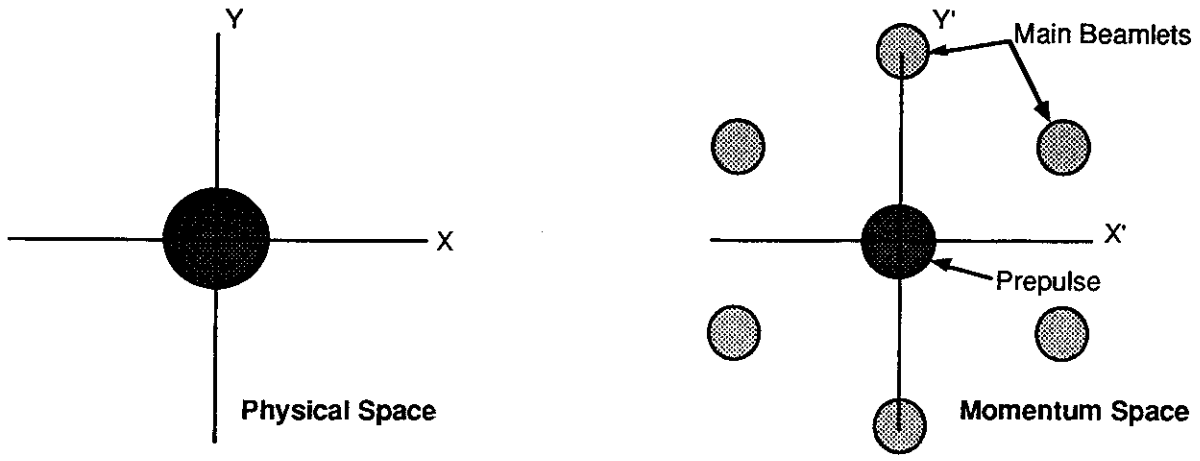

Figure 6.5.2-43. Typical Beamlet Envelopes at Channel Formation Point

The magnetic field associated with the main beam current is sufficient to overcome the initial divergence of the individual beamlets, assuming charge neutralization of the beam within the target chamber (almost certain) and no reverse plasma current (uncertain). This also happens with the precursor, and although the current is down by a large factor, the initial divergence is also much smaller. Figure 6.5.2-44 illustrates the current margin for self-pinching at two possible charge states.

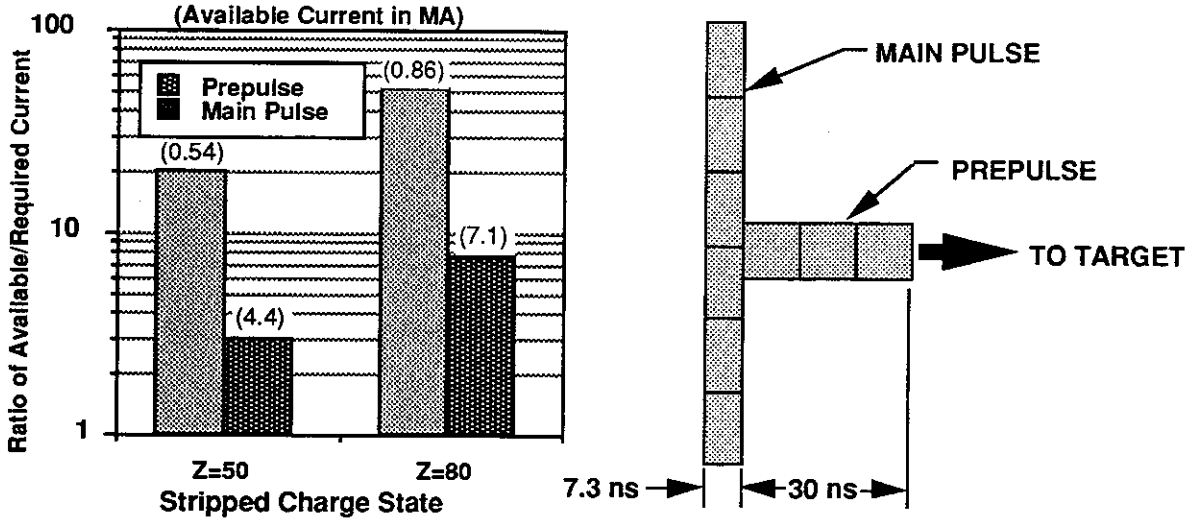

Figure 6.5.2-44. Current Margin for Beam Self-Pinching Indicates Channel Formation Is Feasible in Both Prepulse and Main Beamlets

The equilibrium radius (the radius at which the beam would propagate at fixed size without radial oscillations) is determined by balancing emittance terms ($\sim r^{-2}$) and pinching force ($\sim r^{1}$). Because the initial beam size is much larger than this equilibrium radius, the beam envelope first contracts to a much smaller spot, and then oscillates with period λ_B as it propagates, its maximum radius equaling the injection radius as shown in Figure 6.5.2-45. The two pinched beams (one from each side of the target chamber) continue until they strike the energy converter regions of the indirect drive target.

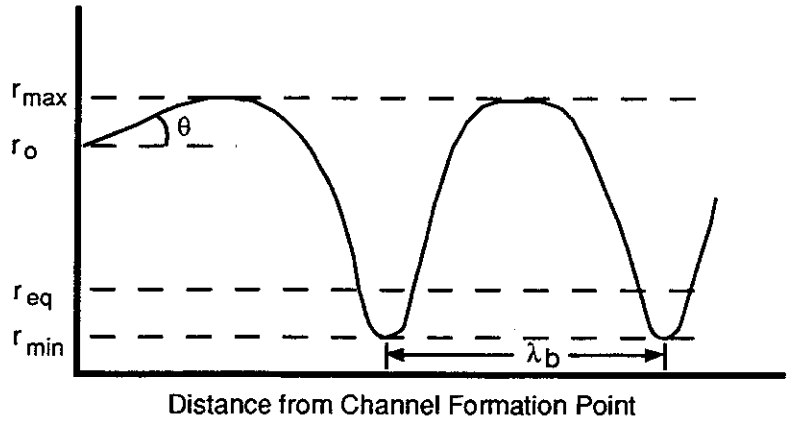

Figure 6.5.2-45. Beam Envelope Oscillation in Channel

Justification. This pinched mode transport is highly speculative, and researchers differ as to whether the plasma return current is sufficiently low that the net current is able to drive the pinch. However, the many engineering advantages of this mode of transport (few reactor blanket penetrations, small shielding penetrations, high chamber operating pressure) make one loathe to give it up without a fight. A definitive experiment would require an ion beam with parameters close to those of the HIF driver.

In the late 1970's, there was considerable interest in the process, but the work was halted due to program redirection. Simulations performed to date have yielded mixed results. Early simulations of the process indicated that the beam head would not be contained. The simulations assumed that the beam entered the gas-filled chamber in a low charge state and that it stripped to a higher charge state and ionized the gas as it went toward the target. As the gas ionized, the conductivity would increase and the magnetic flux would be frozen in, generating a plasma return current to nearly cancel the beam current. The beam head would expand rather than contract.

Recently, preliminary simulations have indicated that the situation can be quite different in a preionized channel with no initial current (as opposed to a pre-formed Z-pinch as suggested for light ion drivers). Such a pre-formed plasma channel can be formed by directing intense pulses of laser light along the chamber axes, ahead of the

precursor beams, and ionizing the gas in the target chamber. In these simulations, the fully stripped ion beam draws electrons ahead of it into it. This sets up an axial current that reinforces the pinch current rather than canceling it, permitting transport to the target.

Other Issues. Assuming that the pinched mode of transport is given, some questions remain concerning the ability to hit the target. The beam must hit a target 5 m away with a high accuracy of the order of ±1 mm.

First note that the ion beam is so stiff that it is virtually impossible to steer the beam to track the target. Instead, the target delivery system must deliver the target to the right place. The "target" for the beams is actually the two conversion plugs in the indirect drive pellet, each of which must face one of the pinched beams.

The questions are whether imbalances between net beamlet current causes steering of the pinched beam on a shot-to-shot basis and whether variations within the beamlets during a pulse can cause steering. The answers depend on whether or not the precursor provides a focusing mechanism; if not, then limits on the allowable current variations in the beamlets can be set for the axisymmetric 6-beam system set on an 8.5-degree cone.

Reference for 6.5.2

1. B. Badger, et al., "HIBALL II: An Improved Heavy Ion Beam Driven Fusion Reactor Study," UWFDM-625, University of Wisconsin Fusion Technology Institute (1984); see also KfK-3480, Kernforschungszentrum Karlsruhe (1984) and FPA-84-4, Fusion Power Associates (1984).

2. Ho, et al., 1987 IEEE Particle Accelerator Conference, 87 CH2387-9, pg. 1994.