Manipulation of High-Current Pulses for Heavy-Ion Fusion

W. M. Sharp, D. A. Callahan, A. Friedman, and D. P. Grote

This paper was prepared for submittal to the
Computational Accelerator Physics Conference
Williamsburg, Virginia
September 24-27, 1996

10/28/96

Lawrence
Livermore
National
Laboratory

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Manipulation of High-Current Pulses for Heavy-Ion Fusion*

W. M. Sharp, D. A. Callahan, A. Friedman, and D. P. Grote

*Lawrence Livermore National Laboratory, University of California
Livermore, California 94550, U. S. A.*

Abstract. For efficient induction-driven heavy-ion fusion, the current profile along a pulse must be modified in a non-selfsimilar manner between the accelerator and the target. In the accelerator, the pulse should have a duration of at least 50 ns in order to make efficient use of the induction cores, and the current should be nearly uniform along the pulse to minimize the aperture. In contrast, the optimal current profile on target consists of a main pulse of about 10 ns preceded by a longer low-current "foot." This pulse-shape manipulation must be carried out at the final pulse energy (5-10 GeV for 200 amu ions) in the presence of a large nonlinear longitudinal space-charge field. A straightforward method is presented here for doing the required pulse shaping. Induction-cell voltages are generated using idealized beam profiles both in the accelerator and on target, and they are verified and checked for error sensitivity using the fluid/envelope code CIRCE.

1. Introduction

For a heavy-ion beam to be an effective inertial-fusion driver, the energy deposited on the target must be carefully controlled. Although the details of energy deposition depend on the target design, a typical indirect-drive target requires a main pulse of about 10 ns duration, with several beams delivering a total of 5-10 MJ. This pulse is normally preceded by a low-current "foot" containing about a tenth of the main-pulse energy, and through the entire pulse, ion velocities should be nearly uniform to minimize chromatic aberration in the final focus. For induction accelerators, which work best with a small number of high-current beams, the peak current on target can be as high as 10 kA per beam, and in the absence of charge neutralization, the longitudinal space-charge fields can be as high as 10 MV/m.

Producing a beam like this at the target is problematic because the optimal beam in an accelerator has a very different current profile. For induction drivers, the aperture is minimized for beams with a uniform line-charge density and at most a small head-to-tail velocity increase to give slow compression. Furthermore, the pulsed-power is used most efficiently when the beam duration is much longer than the pulse rise and fall times, leading to pulses that are ten to twenty times longer in the accelerator than on target. These longer beams also reduce the longitudinal space-charge fields at the beam ends to a manageable magnitude. Considerable manipulation is needed to shape this optimal accelerator pulse into the form needed on target, and the large space-charge field of the final beam prevents this shaping from being done during the final stages of compression, since induction accelerators can presently produce an average gradient of about 1 MV/m.

The conventional scenario for generating the current profile needed on target is to give a pulse carefully controlled current and velocity profiles before it enters a "drift-compression" lattice, where alternating-gradient (AG) focusing is applied, but the beam is not accelerated. The imposed head-to-tail velocity variation, or "tilt," is as large as 5%, and it is "tailored" so that the beam compresses to the desired final profile just as the space-charge field causes longitudinal-velocity "stagnation" at all points along the pulse. This plan requires intricate manipulation of the beam after the main acceleration sequence, probably in a straight "shaping" lattice between the accelerator and the drift-compression section.

To date, the principal published work addressing the question of shaping a constant-current beam into a form suitable for final compression is a paper by Ho, Brandon, and Lee¹. That paper used analysis and one-dimensional numerical simulations to model a beam-compression sequence that produced uniform current and velocity on target. The effects of longitudinal space charge were ignored during shaping, on the grounds that shaping could be done in a sufficiently short time. The present paper presents an alternate procedure for shaping intense heavy-ion pulses that extends the initial work in three major ways. First, we look at the

* Work performed under the auspices of the U. S. Department of Energy by Lawrence Livermore National Laboratory under contract W-7405-ENG-48.

problem of producing a final beam with a prescribed profile, rather than one with uniform current. Second, we include the longitudinal space charge during compression and compensate for it by modifying the shaping fields. Finally, the effect of the changing radial structure of the beam is included approximately in the model. The new procedure is detailed in Section 2 of this paper, and the numerical simulation of a simple case is presented in Section 3. The concluding section gives a brief summary of conclusions.

2. PROCEDURE

The procedure here for calculating the pulse-shaping fields has four main steps, each of which is detailed below. The current and velocity profiles needed at the beginning of the drift-compression lattice are first determined by simulating the target beam drifting backward through the section. A corresponding beam with the same total charge and with user-specified current and velocity profiles is then constructed at the start of the shaping section, and a set of shaping fields to be applied in that lattice is calculated, ignoring the effects of the longitudinal space-charge field. Finally, longitudinal-control fields, referred to here as “ears,” are added to the shaping fields to approximately balance the space-charge field. The combined fields are validated and checked for error sensitivities by numerical simulations using the code CIRCE², which models the transverse beam dynamics using envelope equations and treats longitudinal dynamics using a Lagrangian fluid approximation. The beam is divided lengthwise into $I - 1$ “slices,” each characterized by an enclosed charge δq_i , and the positions and velocities of the slice boundaries are calculated from fluid-like equations. The code has been benchmarked against the three-dimensional particle-in-cell code WARP3d³ and is found to give reliable results, provided that the normalized emittance does not increase significantly and there is little longitudinal mixing.

2.1 Calculate beam profile entering drift section

The current and velocity profiles on target must be selected first. This choice obviously depends on the details of the target, and the target in turn must be designed with limitations of the accelerator and final focus in mind. Here, we use a simple “generic” target beam, and we sidestep problems of the final focus and transport to the target by assuming that there is little change in the current and velocity after final focus. Instead, we set up a beam that is matched to the quadrupole lattice in the drift-compression section.

We would like to obtain the appropriate beam at the beginning of the drift-compression section by numerically running the chosen target beam backward in time through the compression lattice. However, due to limitations of the modeling code CIRCE, we instead run the equivalent problem of the target beam expanding forward in time through the spatially reversed compression lattice. If the scaled longitudinal velocity $\beta \equiv v_z/c$ along the target beam deviates from the charge-density weighted average velocity $\bar{\beta}$, the sign of $\beta - \bar{\beta}$ must also initially be reversed. After this beam is allowed to expand, the current is equated to the desired current I_{bf} at the endpoint z_f of the shaping section, and the desired velocity there β_f is set to $2\bar{\beta} - \beta$. Here, the f subscripts denote final values after beam shaping, and c is the speed of light in vacuum.

2.2 Calculate beam profile entering shaping section

To calculate an appropriate initial distribution, the shaping lattice must first be specified. In most cases, this lattice consists of AG focusing magnets with single or clustered acceleration gaps at one or more locations in a half-lattice period. In principle, the lattice could also include bends and higher-order focusing elements, but these are not considered here. The minimum number of acceleration cells is set by breakdown electric field across a cell and the velocity tilt required at the start of drift compression. If we assume instantaneous acceleration in gaps, the velocity change in the n th gap due to a voltage $V_n(\tau)$ is approximately by

$$\Delta\beta_{i,n}^2 \equiv \beta_{i,n}^2 - \beta_{i,n-1}^2 \approx \frac{2qe}{\gamma_{i,n}^3 M c^2} V_{i,n}, \quad (1)$$

where q is the ion charge state, e is the elementary charge, M is the ion mass, and $V_{i,n} \equiv V_n(\tau_{i,n})$ is the voltage at the time $\tau_{i,n}$ that the i th beam slice traverses the gap. The Lorentz factor here $\gamma_{i,n} \approx [1 - \frac{1}{2}(\beta_{i,n}^2 + \beta_{i,n-1}^2)]^{-1/2}$ complicates the relation between $\Delta\beta_{i,n}$ and $V_{i,n}$, but since it varies by less than

Figure 1. Cartoon of the beam-shaping sequence.

0.4% during shaping and compression, we replace the factor here by the nominal value $\bar{\gamma} \approx (1 - \bar{\beta}^2)^{-1/2}$. According to Eq. (1), the average voltage that must be applied to the i th beam slice in each of the N acceleration modules is

$$\bar{V}_i \approx \frac{1}{\alpha N} (\beta_{i,N}^2 - \beta_{i,0}^2), \quad (2)$$

where $\alpha \equiv 2qe/\bar{\gamma}^3 Mc^2$. The subscript $n = 0$ here denotes values at the start of the shaping section z_0 . We choose to apply the same voltage to the beam head in all cells, and likewise, to the tail, so from Eq. (2), the required head and tail voltages are $V_{head} = \bar{V}_1$ and $V_{tail} = \bar{V}_I$ respectively. We take the maximum electric field that can be held off by the gap insulators to be $E_{max} \approx 10^7$ V/m. Although advanced insulators can hold off twice this voltage for up to 50 ns, the longer beams required for heavy-ion fusion (HIF) and the need to minimize cost lead us to use the more conservative figure. If all acceleration gaps are taken to have the same length L_{gap} , then the minimum number of cells needed in the shaping section is $N_{min} \approx \max(V_{head}, V_{tail})/(E_{max} L_{gap})$.

Due to the low average gradient in induction accelerators, the shaping section must typically be longer than 100 m, so the beam duration is expected to decrease measurably as the beam traverses the section. To construct a suitable beam profile at z_f to be manipulated into the calculated profile at z_f , we first estimate this change in beam duration, again assuming that the beam head and tail respectively see V_{head} and V_{tail} in each cell. In the cartoon of Fig. 1, this calculation corresponds to finding the trajectories $\tau_{1,n}$ and $\tau_{I,n}$ of the head and tail slices back to z_0 . The transit time between cells in the absence of longitudinal space charge is given by

$$\Delta\tau_{i,n} \equiv \tau_{i,n+1} - \tau_{i,n} \approx \frac{\Delta z_n}{c\beta_{i,n}}, \quad (3)$$

where $\Delta z_n \equiv z_{n+1} - z_n$ is the distance between the centers of the n th and $(n+1)$ th gap, and from Eq. (1),

$$\beta_{i,n} = \left(\beta_{i,0}^2 + \alpha \sum_{m=1}^n V_{i,m} \right)^{1/2}. \quad (4)$$

Defining z_{N+1} to be z_f and $\tau_{i,N+1}$ to be τ_i there, we iterate Eqs. (3) and (4) for the head and tail slices backward to z_0 , obtaining $\tau_{i,n}$ and $\beta_{i,n}$ along those trajectories. The desired current profile $I_{b0}(\tau)$ at z_0 is then scaled to have the appropriate duration $\tau_{I,0} - \tau_{1,0}$ and the correct total charge $Q = \sum_{i=1}^I \delta q_i$.

To complete the specification of the initial beam, we must choose times $\tau_{i,0}$ for $1 < i < I$ such that each slice contains the same current δq_i as the beam at z_f . These intermediate τ values are found numerically

by requiring that the τ integral of I_{b0} from $\tau_{1,0}$ to $\tau_{i,0}$ equals the charge in the same section of beam $Q_i = \sum_{j=1}^i \delta q_j$.

2.3 Calculate pulse-shaping fields

The values of $V_{i,n}$ and $\tau_{i,n}$ now must be specified for $1 < i < I$ and $1 < n < N$. We would like $V_{i,n}$ to be close to $\bar{V}_i$ so that all cells have to generate approximately the same voltage, but taking $V_{i,n} = \bar{V}_i$ does not, in general, give the correct slice arrival times at z_f . Instead, we write $V_{i,n} = \bar{V}_i + \delta V_{i,n}$ and choose the voltage difference $\delta V_{i,n}$ to satisfy two constraints:

$$\sum_{n=1}^N \delta V_{i,n} = 0, \quad (5a)$$

$$\frac{1}{c} \sum_{n=1}^{N-1} \frac{\Delta z_n}{(\beta_{i,0}^2 + n\alpha\bar{V}_i + \alpha \sum_{m=1}^n \delta V_{i,m})^{1/2}} = \tau_{i,N} - \tau_{i,1}. \quad (5b)$$

The first constraint guarantees that the final slice velocity is unchanged, and the second, from Eq. (3), ensures that the total transit time from z_1 to z_N is correct. We simplify Eq. (5b) by assuming that the contribution of the voltage perturbation to $\beta_{i,n}$ is small, so that the expression can be linearized in $\delta V_{i,n}$:

$$-\frac{\alpha}{2c} \sum_{n=1}^{N-1} \frac{\Delta z_n}{(\beta_{i,0}^2 + n\alpha\bar{V}_i)^{3/2}} \sum_{m=1}^n \delta V_{i,m} \approx \tau_{i,N} - \tau_{i,1}^0. \quad (6)$$

Here, $\tau_{i,n}^0$ denotes the slice arrival times along constant-voltage trajectories, given explicitly by

$$\tau_{i,n}^0 \equiv \tau_{i,1} + \frac{1}{c} \sum_{m=1}^{n-1} \frac{\Delta z_m}{(\beta_{i,0}^2 + m\alpha\bar{V}_i)^{1/2}}. \quad (7)$$

As a sample solution of these equations, we take $\delta V_{i,n}$ to vary linearly with n , so that

$$\delta V_{i,n} = A_i \left(\frac{n-1}{N-1} - \frac{1}{2} \right). \quad (8)$$

This expression satisfies Eq. (5a) for any value of A_i , and an appropriate A_i approximately satisfying Eq. (5b) is found by substituting Eq. (8) into Eq. (6). We then calculate the voltage adjustments $\delta V_{i,n}$ from Eq. (8) and find the corresponding $\tau_{i,n}$ values by iterating Eq. (3), using the new voltages $\bar{V}_i + \delta V_{i,n}$ in the $\beta_{i,n}$ expression of Eq. (4). The resulting shaping voltages and time data are then written to an external file for use in CIRCE. This procedure for calculating the shaping voltages is used in the sample problem of Section 3, but it is in no sense optimum. Performance measures for global optimization can readily be formulated, but the large number of shaping cells and beam slices typically used makes the problem unwieldy. Various optimization schemes are under study.

The pulse-shaping method here is predicated on the assumption of instantaneous acceleration in the induction cells. We have compared the results of this approximation with an analytic expression obtained for a finite-length gap, and we find that the relative error in $\tau_{i,n}$ due to the approximation is, to lowest order,

$$\frac{L_{gap}}{\Delta z_n} \left[\frac{\Delta \beta_{i,n}^2}{4\beta_{i,n}^2} \right]^2, \quad (9)$$

where $\bar{\beta}_{i,n}c$ is the average velocity of the i th slice going through the n th cell. Even when the gap length is a significant fraction of Δz_n , this error is still small due to the low fractional energy change in a single induction cell of a heavy-ion driver. For example, the factor $\Delta \beta_{i,n}^2 / \bar{\beta}_{i,n}^2$ is approximately 10^{-5} for a 10 GeV

Figure 2. The current and velocity profiles of an idealized heavy-ion beam (a) after acceleration, (b) after shaping, and (c) at the target. Profiles are plotted as functions of arrival time τ relative to the beam head.

driver with conventional acceleration modules. We conclude that instantaneous acceleration is a very good model for the cases under consideration.

2.4 Calculate ear fields

The straight-line trajectories assumed in calculating the shaping fields are, of course, only valid in the absence of longitudinal space-charge fields. To calculate appropriate ear fields, we use CIRCE to simulate the beam transit through the shaping section, applying the shaping fields, but with the longitudinal component of the space-charge field artificially set to zero. When the beam enters each shaping cell, the beam profile is used to calculate the space-charge field at that point, and this field is then used to construct the appropriate ear field there.⁴ This ear field is added to the shaping field, and the combined field is written along with timing data to an external file for use on later runs.

3. SAMPLE CASE

The shaping procedure is illustrated here using an idealized analytic beam-current profile on target and a simple lattice for shaping and compression. The purpose is primarily didactic, since no effort has been made either to optimize the lattice or to choose a target profile appropriate for a specific target design.

3.1 Target-beam characteristics

The beam on target is deliberately chosen to stress the pulse-shaping algorithm. An ion mass of 200 amu with a final energy of 10 GeV is assumed to give the beam high rigidity, and the 10 kA peak current on target, which is appropriate when there are four beams there, exaggerates the effect of space charge. The final normalized x and y emittances are taken to be 10^{-5} m-rad, which should allow the beams to be focused to a 3 mm radius spot on the target. The functional form of the current has the general features needed for a target with two end plugs, similar to that specified by Ho⁵, but the actual analytic form has been chosen for its simplicity and convenience.

It is found that the current profile on target cannot be specified arbitrarily. For beams in which all ions have the same mass and momentum, several classes of profiles are unusable due to pathological dynamics when they expand. For example, the expansion wave from the leading edge of the main pulse tends to form a steep density gradient as it expands into the lower-density foot, provided that the space-charge force is

Figure 3. The (a) shaping fields and (b) ear fields at selected acceleration modules in the shaping section. No optimization has been done to smooth ear fields.

sufficiently high. If “wave-breaking” occurs, ions in the expanding wave overlap the slower ions in the foot, producing a multi-valued velocity distribution that is impossible to produce using external fields. Similarly, if the target distribution has two peaks separated by a lower current region, a beam distribution with two distinct velocity components in the same spatial region is formed if the two expansion waves meet. In the fluid-like longitudinal dynamics model of CIRCE, where the charge in every beam slice is constant, the density develops nearly singular spikes at the points where a multi-valued distribution would form in reality. Consequently, we disqualify a target profile when such spikes form during expansion.

The tendency to form a multi-valued distribution is reduced by increasing the current in the foot, which increases the velocity of electrostatic waves there, or by lengthening the rise time of the main pulse, which reduces the space-charge force driving the main-pulse expansion. A third approach is to give the main pulse a slightly higher velocity, so that it “catches up” with the foot as the beam compresses. All three strategies have been used to generate the target profile shown in Fig. 2c. Here, the 2 kA foot current is about a factor of two higher than would be optimal, but choosing the foot velocity to be about 1% lower than in the main pulse partially compensates for the higher current by increasing the focal spot size, so that some current misses the capsule end plugs. Although this profile expands in a reversible fashion, the deliberate defocusing of the foot wastes energy and lowers the capsule gain. Careful optimization is clearly needed here.

3.2 Beam characteristics before and after shaping

The 45-period drift-compression section used here is a simple AG lattice with a half period of 4.18 m, with the quadrupole strength and occupancy chosen to give an maximum undepressed phase advance of 75° . When the target-beam in Fig. 2c is allowed to propagate backwards without longitudinal confinement through this lattice, it develops the current and velocity profiles shown in Fig. 2b. Here, as in Fig. 2c, we have reversed the velocity so the plot is appropriate for a compressing beam. For this case, the beam duration has approximately quadrupled to 200 ns, and the rarefaction waves from the ends of the main pulse have met, leaving the velocity roughly linear. We also note that the beam tail has spread out more than the head during expansion because the short fall time of the main pulse leads to a large space-charge field there.

The 60-period shaping section for this case is also a simple AG lattice with the same nominal phase advance and 4.18 m half period as the drift-compression section. However, the beam-pipe radius is smaller, reflecting the lower line-charge density, and most of the space between quadrupoles is filled with acceleration gaps. A cluster of ten 10 cm gaps is inserted between each quadrupole doublet, allowing an average gradient in excess of 1 MV/m. We assume that the beam from the accelerator has a uniform longitudinal velocity,

Figure 4. Comparison of the (a) current and (b) velocity of the simulated beam (wide line) and assumed beam (narrow line) at the target.

although allowing a small tilt would reduce the velocity change that must be imposed in the shaping section. When the beam ends are followed backward through this lattice, we find that the beam duration increases to about 300 ns at the starting point. Choosing the rise and fall times to be 25% of the beam duration, we can then set up the initial beam profile shown in Fig. 2a. We see by comparing Figs. 2a and 2b that a head-to-tail velocity tilt in excess of 4% must be imposed by the shaping section.

3.3 Shaping and ear fields

A selection of waveforms for this illustrative case, taken at five equally spaced modules, is shown in Fig. 3a. Here, the duration of the waveform is seen to decrease at successive z positions as the beam compresses. The procedure in Section 2c generates acceptably smooth shaping voltages, but the corresponding ear voltages, shown in Fig. 3b, are seen to have considerable high-frequency content. This feature is a consequence of the shaping voltages being adjusted along each slice trajectory, without regard for voltages along neighboring trajectories. An erratic current profile results when $\tau_{i,n}$ at any z_n are irregularly spaced, producing longitudinal fluctuations in the space-charge field there. This irregularity in the ear fields would be reduced by a suitable optimization procedure.

3.4 Dynamics

When the idealized accelerator beam in Fig. 2a is transported through the shaping and drift-compression sections with ears and shaping fields applied, the final beam has the basic features of the intended target beam, as shown in Fig. 4. The wide lines in the figure give the current and velocity profiles of this test pulse, and the narrow lines are the intended profiles from Fig. 2c. The greatest departures from the intended profiles occur at the tail of the main pulse, where the longitudinal space-charge field is largest. Presumably, these departures occur because the high space-charge fields make the beam particularly sensitive to errors in the current or velocity at the start of the drift-compression section. The ear fields are found to be important to replicating details of the intended current profile. When the ears in Fig. 3b are not added to the shaping fields, the final current profile of the main pulse assumes a nearly triangular shape, and the rising and falling portions stretch out by a total of 13 ns.

A preliminary study of the sensitivity of the final profiles to errors in the shaping fields has been made. It is found that random 2.5% rms errors in the magnitudes of the shaping fields produce final profiles that are not significantly worse than the error-free case in Fig. 4, but progressively larger discrepancies from the intended profiles are seen for random 5% and 10% errors. Systematic errors of the same magnitude cause much worse discrepancies in each case and are probably unacceptable even at the 2.5% level. Random

shaping-field timing errors with an rms magnitude of 10 ns are found to cause a head-to-tail current increase exceeding 10% across the tops of both the main pulse and the foot. This deviation from the optimum current profile is probably acceptable in practice, but for a random 25 ns timing jitter, the beam no longer compresses successfully, due to the formation of large density spikes near the beam tail. These sensitivities illustrate the need for high-precision pulsed power in the shaping section.

A much more thorough examination of errors sensitivities is obviously needed. In particular, the effects of limiting the bandwidth of the shaping voltage should be checked. Also, these cases should be simulated using WARP3d to verify that the fluid-like longitudinal dynamics of CIRCE has not introduced any unphysical effects here.

4. SUMMARY

The work presented here suggests that an ion pulse from a heavy-ion driver can be systematically shaped by a series of small changes. The procedure here has not been optimized or coordinated with target simulations, but it nonetheless illustrates that predetermined beam profiles can in principal be produced. Generating the required pulses with adequate precision may prove challenging because existing pulse-generation technology imposes limits on both the magnitude and frequency content of the shaping signal. A more fundamental problem is the limited choice of target profiles that can be formed by drift-compression in vacuum. Various ways to circumvent this problem as being considered, such as the partially neutralized transport, the application of "mid-course corrections" during compression, and the use of separate beams, possibly with different ion masses, for the foot and main pulse.

5. REFERENCES

- ¹ D. D.-M. Ho, S. T. Brandon, and E. P. Lee, *Part. Accel.* **35**, 15-42 (1991).
- ² W. M. Sharp, J. J. Barnard, D. P. Grote, S. M. Lund, and S. S. Yu, "Envelope Model of Beam Transport in ILSE" in *Proceedings of the 1993 Conf. on Computation Accelerator Physics*, Pleasanton, CA, February 1993, pp. 540-548.
- ³ A. Friedman, D. P. Grote, and I. Haber, "Three Dimensional Particle Simulation of Heavy Ion Beams," *Phys. Fluids B* **4**, 2203-2210 (1992).
- ⁴ W. M. Sharp, D. A. Callahan, and D. P. Grote, "Longitudinal Dynamics and Stability in Beams for Heavy-Ion Fusion," in *Proceedings of the 1995 Heavy-Ion Fusion Symposium* Princeton, NJ, 6-9 September 1995, in press.
- ⁵ D. D.-M. Ho, private communication.