

Phase-Space Considerations

Roger Bangerter

Historically phase-space considerations have been among the factors that determine the parameters and characteristics of accelerators for fusion. Targets and focusing systems constrain all six dimensions of phase space and therefore place limits on both the transverse and longitudinal emittance of the beam(s) emerging from the accelerator. Although it is, in principle, possible to reduce emittance, the techniques for doing so are difficult to apply to inertial fusion. In general, emittance is expected to increase as the beam is accelerated.

The requirements for HEDP are similar to those of fusion. In both cases the focussed intensity and the specific energy deposition are the fundamental quantities; but, in fusion, the requirement on the ρ - r product of the fuel sets a minimum scale size for total energy that is not necessarily applicable to HEDP. The quantities that are important for HEDP and fusion are also important for WDMP but the requirements for WDMP are significantly less stringent. Nevertheless, phase-space considerations are also of fundamental importance for WDMP.

There are a number of effects that can, under some circumstances, produce emittance or lead to its growth. For longitudinal emittance these effects include:

1. The longitudinal “temperature” of the source.
2. Coupling of transverse temperature to the longitudinal direction. This coupling is expected to be rapid. The longitudinal temperature rises to a value that is comparable to, but less than, the transverse temperature. In most accelerators that have been designed, the beam is compressed transversely after it leaves the source so the transverse temperature of the beam can be higher than the source temperature, even without transverse emittance growth.
3. Scattering and charge exchange in the source.
4. Longitudinal energy variations or wiggles on the beam emerging from the source. Although there is a 1-D analytic extraction waveform that gives a monoenergetic beam, these wiggles have always been seen in experiments and in simulations, for example, in the simulations shown by Enrique Henestrosa at this workshop. It is not known if a perfect waveform exists in 3-D.
5. Electrons. Fields produced by electrons are often unpredictable. They can produce nonlinear fields that lead to effective emittance growth.
6. Waveform errors in the injector and accelerator -- and in any final ramp used to compress the beam longitudinally.
7. Nonlinear or nonuniform self fields.
8. Mismatches between self and applied fields.
9. Instabilities, waves, etc. in the injector, accelerator, and final compression section.
10. Virtual anode effects.
11. Inductive effects.

12. Rapid transitions between accelerator sections. Steve Lund has recently shown that rapid transitions can have a significant effect on transverse phase space. It is likely that there will be longitudinal analogs of these effects.

These effects can be quite constraining. Consider, for example, some recent modular-solenoid driver cases presented at the 2004 HIF Symposium. These systems were designed to deliver about 6.7 MJ of energy -- appropriate for the hybrid target. Economic considerations limit the number of modules. The cases presented at the Symposium used about 24 modules. Each module accelerated two pulses of neon ions to about 210 MeV. The initial duration of each pulse was approximately 20 ns. There are important technical issues associated with double pulsing; but, for now, assume that double pulsing works and that the bunches can be combined longitudinally as they approach the target. Thus there are 12 beam lines on each side of the target.

Consider first a single beam line. MRC has simulated a 1-beam compression and focusing system based on neutralized drift compression with solenoidal focusing into an adiabatic plasma lens, followed by a 5 mm, 50 kA channel. The results were presented at the VNL PAC meeting in August 2004. The 5-mm channel is appropriate for the hybrid target. Other parameters are an initial beam current of 3.35 kA of singly charged neon, an initial beam radius of 10 cm, a normalized transverse emittance of 8 mm·mr, 147 kJ of beam energy, and an initial pulse duration of 210 ns. These numbers give an average kinetic energy of approximately 210 MeV, an ion velocity of 0.15 c, a line charge density of 75 C/m, and a charge per pulse of 0.7 mC. A total of 24, 147-kJ beams would give 3.5 MJ, so the simulations rigorously apply to only one of the two pulses in each beam line (with a few percent safety factor). The simulations show a remarkably large tolerance to energy spread, a total of 20% head-to-tail corresponding to a half width of 21 MeV. With this energy spread, slightly over 90% of the beam falls within a 5-mm spot at the end of the channel. The duration of the final pulse for a 7-MJ target is of the order of 10 ns, but it must have a special shape and a rapid fall time to avoid wasting energy. It appears reasonable to believe that one could synthesize the pulse by appropriate timing of a number of individual pulses having rounded pulse shapes with a width of 5 ns or a half width of 2.5 ns. Combining 2.5 ns and 21 MeV, we get a longitudinal admittance of 0.0525 eV·s.

For comparison, consider the longitudinal emittance produced by a typical ion source. Simulation numbers are available for the multi-aperture ion source recently designed by Joe Kwan and Dave Grote. The simulated performance of this source is comparable in brightness to the measured properties of typical thermionic sources for potassium. For lighter ions such as neon (or sodium), the multi-aperture sources will likely be superior.

Grote's simulations give 0.57 A of 40-amu ions at 1 MeV. After beamlet merging, the mean radius of the elliptical beam is 1.55 cm and the normalized emittance is 0.9 mm·mr. One can estimate the transverse beam temperature using the familiar formula,

$$\sigma = m (0.5 \beta a)^2 \quad (1)$$

where ϵ is temperature (eV), m is ion mass (eV), and ϵ and a are respectively normalized emittance and beam radius in consistent units (m•rad and m). Using Grote's numbers one obtains $\epsilon = 30.6$ eV. If one scales these sources to different current by changing the number of beamlets, one would expect the temperature to remain invariant since the electrostatic energy of a single beamlet is invariant. Similarly, the temperature should not depend strongly on ion mass because the line-charge density doesn't depend on ion mass.

By what factor must the source be scaled? Obtaining the required 0.7 mC of charge in 20 μ s requires a current of 35 A. For neon, the source would produce 1.4 times as much current as for potassium, or 0.806 A. The source must therefore be scaled up a factor of $35/0.806 = 43.4$ in transverse area or 6.59 in radius. Remarkably, this scaling factor gives a beam radius of 10.2 cm which is essentially consistent with the radius chosen by MRC. (The line-charge density and therefore the radius are more or less invariant in a solenoidal system. Note also that the normalized transverse emittance given by equation (1) would be 8.3 mm•mr.) Anyway, no additional expansion or compression of the beam is needed so 30.6 eV remains a good estimate of the transverse temperature. Assume that one half this temperature is transferred to the longitudinal direction. Transformed to the laboratory, a longitudinal beam temperature ϵ gives a longitudinal energy spread $\Delta T = \pm 2(2\epsilon T)^{1/2}$ where T is kinetic energy. The initial factor of two comes from multiplying the rms velocity spread by 2 to convert to a value suitable for calculating 'edge' emittance. Using the source voltage of 1 MV, one obtains $\Delta T = \pm 11.1$ keV. Multiplying by one half the initial pulse duration (10 μ s) to get longitudinal emittance, one obtains 0.11 eV•s which is slightly more than twice the admittance of the focusing system. Note that this emittance comes only from effect 2 in the list given above.

Now consider double pulsing. If one could combine two pulses perfectly in the longitudinal direction, the longitudinal emittance would double. If the interpulse time is equal to the pulse duration (reset time equal to pulse duration) the longitudinal emittance would increase by a factor of 3 – actually more than a factor of three because of non-negligible rise and fall times. Such a strategy would already have a significant effect on accelerator efficiency because the core losses during reset would equal the core losses during one of the pulses. With a single pulse, the reset losses are much lower because the reset is done slowly.

In summary, based on a consideration of only one of the many sources of longitudinal emittance, and on a specific type of ion source and final focusing system, the estimate just described fails to give an acceptable answer by at least a factor of 4.

Is the estimate just described likely to be optimistic or pessimistic? By working on target design, one might be able to use a rounded pulse shape that is natural to produce. Maybe one could hope to increase 2.5 ns to nearly 5 ns. Or perhaps less than half the transverse temperature might be transferred to the longitudinal direction. On the other hand, the current calculations appear to be optimistic in a number of ways. The neglect of all sources of longitudinal emittance except the second has already been mentioned. It appears to this author, that some of these effects could easily give contributions to

longitudinal emittance that are comparable to or larger than the effect considered. Also, the MRC simulations may be quite optimistic. The initial beam size is assumed to be 10 cm; but, according to equation (2) of Lee and Briggs (LBNL-40774), the rms on-axis solenoidal field would have to be 10.6 T to transport the assumed 75 $\mu\text{C}/\text{m}$. Therefore the peak field would have to be greater than 10.6 T, particularly if one considers the substantial gaps and dead-space needed for a high-gradient machine. This problem is exacerbated if one includes the relatively large rise and fall times needed to keep the ‘ear’ fields within acceptable bounds. In the ‘all-ears’ limit (a parabolic pulse) one would have to either increase the peak line-charge density by 50% or increase the pulse length by 50%. Also, since only one of the two pulse was considered, the line-charge density after compression would have to be increased by an additional factor of two to make the simulations realistic for 24 beams. Furthermore, the simulations assumed unphysical magnetic fields for the electron barrier and they assumed a perfectly linear adiabatic lens and a perfectly linear channel. Effects of stripping and recombination were ignored except at the beginning of the adiabatic lens where stripping to +10 was assumed to be instantaneous and complete. Moreover, the longitudinal admittance of such a system decreases as transverse emittance increases. The initial transverse emittance of the example ion source already exceeds the emittance assumed in the simulation but we have assumed that some of the transverse emittance is coupled to the longitudinal direction. Regarding transverse emittance, it would be necessary to combine 12 beams transversely into a single channel. Even if the distance between beam centers were only 3 beam radii, the transverse emittance in the channel would have to increase by a factor of about 5 relative to the 1-beam case. Finally, there is one effect that remains largely unstudied. At the target, it is necessary to provide return-current channels to carry the current away from the target. These would presumably be perpendicular to the beam channels. The angles of the ions relative to the channel axis are large enough (of the order of 100 mrad) that a small distortion or expansion of the beam channels by the target and/or the return-current channels (for example, a 1 mm expansion 1 cm from the target) could cause significant beam loss. Parenthetically, channel expansion is another effect that should be investigated. In an MHD model, the transverse beam pressure greatly exceeds the magnetic pressure.

There are also a number of effects that must be considered regarding transverse emittance. These include:

1. The temperature of the source. Presumably some of this temperature could be transferred to the longitudinal direction if the longitudinal temperature is sufficiently low. The transfer of longitudinal energy to the transverse direction is expected to be slow.
2. Mismatches and transitions.
3. Errors, nonlinearities, and imperfections in the focusing elements.
4. Instabilities
5. Electrons
6. Misalignments

As in the longitudinal case, these considerations can place important limits on

performance. For example, one can easily show that the normalized transverse admittance of the 5-mm, 50-kA channel described above is approximately 75 mm•mr. If the transverse emittance were to increase by even a factor of two in the machine, to say 15 mm•mr, 12 combined beams would completely fill the transverse admittance of the channel, reducing the longitudinal admittance, in this approximation, to zero. In summary, it appears that the example case fails to satisfy the phase-space constraints – probably by a substantial factor.

What can be done to improve the situation? As noted above, improvements in targets could help. In addition there appears to be some room for design improvements and further optimization of the compression and focussing system. One could also hope for brighter ion sources, although this strategy would work only if the source is the dominant source of emittance. One might also try to use a different final focusing system where transverse combination into a single channel is not required.

Consider a system that utilizes neutralized ballistic focusing. If one considers only transverse emittance and chromatic aberrations, one can easily derive the usual equation for the minimal focal spot radius $r = (2\epsilon/f\Delta)^{1/2}$ where ϵ is a chromatic factor ($\epsilon = 1$ for a simple solenoidal system) and f and Δ are respectively the standoff between the final lens and the target and the fractional energy spread (half width). If one puts in the estimates of emittance and energy spread obtained from the example case, the focal spot is too large by a substantial factor if f is greater than about 1 m.

Of course, one could design a modular-solenoid system that meets the phase space requirements by increasing the kinetic energy and ion mass and/or using more, smaller beams (modules). The question is whether such a system can have reasonable economics and efficiency.

As noted above, the requirements are relaxed significantly for WDMP. Some preliminary estimates at the Workshop suggested that phase-space volume, based on only a few of the considerations listed above, might be okay, but not by a large margin, for some of the WDMP systems. More work is required to do the calculations accurately and to include all the important effects.

In conclusion, phase-space considerations are important. Many accelerator options for WDMP and HEDP have been suggested but only a very small number can be pursued. It is likely that careful consideration of phase-space requirements will eliminate a number of the options – particularly if the ability to extrapolate from WDMP to HEDP and/or fusion is believed to be an important consideration.