Induction Accelerator Buncher for Storage Rings

Andy Faltens

A separate induction accelerator buncher following after the storage rings is one of the options for final bunch compression. The other option is to apply the bunching voltage within the ring, but this requires a low-frequency, high gradient accelerating structure within the ring and a large aperture, which are difficult to do and undesirable. The induction accelerator buncher option here differs from the bunching function in a standard induction accelerator scenario in that here, a separate buncher is required, whereas in the induction linac the bunching function is accomplished by ramping the acceleration voltages near the end of the machine. This is a minor difference, but one that allows consideration of a bipolar buncher, which has no net acceleration. The other major difference is that the currents per beam to be bunched are smaller than in the straight induction linac, permitting use of transversely smaller, and hence less expensive, structures.

Beam Transport

A simplified version of the beam transport equations for a space-charge dominated beam in a quadrupole channel, useful for a rough layout of the buncher geometry is:

$$I \approx 8 \times 10^5 \beta^2 \eta B a$$

where a is the beam radius, B is the field at the beam radius, η is the fractional occupancy of the quadrupole field in the half-period of length L, and β is the relativistic factor.

The focusing field decreases linearly towards the center from its value at the superconducting windings, at radius rw. With niobium-titanium superconductor, there is a JB operating curve relating field at the superconductor and the current density. For our purposes, we take this to have an optimum of $6.5 \times 10^4 \, A/cm^2$ at a field of 6 T, with the current density being the macroscopic current density, including copper, electrical insulation, and the empty space within the windings. Figure 1 shows a superconducting quadrupole with windings and compression collars for an isolated channel. Inside the radius of the winding space must be allocated for the thermal insulation and vacuum chamber. In a large array, the adjacent channels can be positioned to aid in the fields of any channel and effectively halve the required amount of superconductor and collars. In a large array the radius of the windings for each channel has an optimum in the sense of transporting the most current averaged over the array. If the beam aperture is very small, then beam clearance, thermal insulation, the vacuum chamber and conductor compression collars take up a large fraction of the quadrupole, with the result that the field at the beam edge is low. If the beam aperture is very large, then the field at the beam will be nearly equal in value to the field at the windings, but the gradient of the field, $dB/dr \cong B/r$ will be low, and the transportable current density, J, which is proportional to the gradient will be low. From considerations such as these, an optimum channel would have a=6 cm and the pitch of the unit lattice cell would be about 24 cm. From the transport equation, the current that could go through this lattice is 2.5 kA per channel, which is 80 times more than required at the beginning of buncher. Advantage can be taken of this by using a lattice with a very low occupancy and a long half period, thereby reducing the transportable current. However, it is not desirable to have very long stretches of accelerating column without some means of impeding electron backflow, so we limit the half-period length to, say, 5 m and 5 MV per section. This is shown in Figure 2. Having chosen the maximum half period, and needing to transport a given amount of current, the channel and the quadrupole dimensions are determined. For the low intensity beams in the buncher region the channels could be made smaller than the optimum and the transverse pitch could be reduced to about one half of the optimum pitch, or 12 cm.

The lattice cell size and the number of beams determine the quadrupole cluster size. For high field quadrupoles, additional windings are required around the periphery of the cluster to terminate the fields in a way that makes each channel appear as part of an infinite array. This solution requires no iron, but increases the array diameter by one lattice spacing (Figure 3). At low fields, less than 1.8 T, an iron boundary can be used to terminate the array and to isolate channels from each other, thereby facilitating inclusion of correction windings, and saving a little space radially. For a large array such as contemplated here the difference is minor.

The size of the array of beams determines the diameter of the accelerating column, and this in turn determines the size of the inner diameter of the induction cores which surround it. The induction cores are usually enclosed by metal conductors, except for an azimuthal gap at the outer radius, and insulated with transformer oil. At the short pulse lengths required for bunching, the radial buildup of the induction cores is small, so to a fair level of approximation the size of the ID also determines the minimum filling time for electrostatic energy within the structure, and hence the useable risetimes for it. The risetime is essentially the time required for electromagnetic waves to travel around the periphery of the core-enclosing metal structures. The magnetic material characteristics are a minor complication compared to the electromagnetic problem. This result, in the lumped element approximation can be interpreted as charging up a gap capacity through the pulse forming network impedance. What matters is how closely the drive line is matched to the characteristic impedance of the structure. For efficiency reasons, the drive impedance is usually chosen to match the beam current and the induction core current, but the risetime or transient response is largely a function of the metal structure. To some extent this structure can be modified to take transient response into account. If the minimum spaces are used everywhere, in an attempt to get maximum acceleration rates, then the geometry is based on breakdown considerations, and the result will be a low impedance or high capacitance structure which takes a longer time to energize.

Figure 4 shows a typical induction module. The "gap" which is charged up is not the vacuum gap through which the beams pass but the dielectrically loaded insulating gap which surrounds the core. The electric field distribution around the core changes with time, as the core saturates, so describing that process in terms of a capacitance is a simplification. The vacuum gap itself, through which the beams pass, is smaller in size and the speed of electromagnetic waves within it is the speed of light in vacuum, so it can be treated as having a constant voltage at any instant in time. For a structure of 40 cm radius with an applied voltage ramp of 100 ns duration this is a good approximation.

Option for Bipolar Pulsing

Bipolar pulsing is a possible option that could in principle double the longitudinal focusing field gradient, from approximately 1 MV/m to 2 MV/m, and reduce the amount of core material required, but at the cost of a more elaborate pulser. One solution is to use a separate pulser for each polarity of the pulse. At a considerably longer pulse duration this has already been done in some induction linac double pulsing experiments, where one pulser was used for generating the acceleration pulse and a second pulser was used to provide a powerful reset pulse, but essentially all induction linacs have been unipolar. Developmental tests would have to be made to see if the time separation between the pulses could be reduced to the 30 ns to 50 ns region. The use of saturating iron core reactors should also be considered in this application to isolate the pulsers from each other, because there is a difficulty in getting semiconductor devices to function quickly enough. Also, the accelerating column for a bipolar pulse would be somewhat longer than for a unipolar pulse. Except for mentioning the bipolar pulse possibility, this report is confined to the more conservative unipolar case.

Pulsers and Cores

For the most recent 24-beam array and the 12 cm assumed transverse lattice period, the required acceleration column inner diameter is 70 cm. Allowing a few cm for the thickness of the column, the core inner diameter would be 80 cm. The economical core materials have a usable flux swing of 2.5 T, which must be derated by the radial and axial packing fractions to 1.6 T to take account of the required insulation. From

$$\int V dt \approx (\Delta B)(\ell)(\Delta r)$$
$$(0.5)(10^6)(10^{-7})(2) \approx (1.6)(1)(0.06)$$

we see that a core buildup of about 6 cm is all that is required, on the average, so a peak buildup of 10 cm is adequate. The factor of 0.5 in the volt second integral arises from the required bunching waveform being a ramp; an equal amount has been added to account for the extra core needed to provide a high impedance to the pulser at the end of the pulse.

The average diameter of the cores is 90 cm, from which the circumferential travel time around the structure is about 14 ns in a dielectric medium with refractive index 1.5. If the drive line is well matched to the structure impedance, then this time is the total fill risetime; if there is a substantial mismatch, then several reflections are required, and this time could increase by a factor of a few. With the voltage pulse applied at a single feed point, the longitudinal electric field within the beam cluster will rise almost in unison, with a time delay between the terminal voltage and the gap voltage of about 7 ns. The large physical size of the structure effectively acts as a low pass filter between the applied voltages and what the beam sees. A similar low pass action also occurs from the transit time factors of the individual beams passing through the acceleration gap. The electrical geometry is that the voltage from any one pulser is applied across a small circumferential gap at the outer diameter of a large diameter pipe, and the fields spill out to a distance of about one radius on either side of gap, which means that an ion with a β of 0.32 will pass through its fields in something like 6 ns. Taken together, the induction cell structure and accelerating gap geometry limit the risetime to 20 ns, which is adequately far from the desired voltage risetime at the exit of the buncher to be a serious concern.

The actual risetime from a hydrogen thyratron switch will most likely be the slowest constituent of the pulse, in the 50-100 ns region. At the entrance to the buncher, it would have to be slowed down further with a small, lumped element compensation network, to match the desired linear risetime. At the exit end of the buncher, if required, the pulser output waveform could be slightly steepened with a saturating reactor and then linearized. Each economical thyratron pulser is capable of about 100 MW of pulse power, which we approximate as a 15 kV voltage with 6 kA available at the induction module.

The main requirements at the induction module arise from the beam current and the core magnetizing current. The bunching scenario is still subject for further optimization, and the required ramp and beam current change somewhat over the length of the buncher, so a small selection of module and waveform types will be required. Here we assume as a specific representative example a beamlet current of 50 amperes per channel and a linear voltage ramp of 100 ns, followed by a decay of another 100 ns. The total beam current thus is 1.2 kA, constant over the pulse duration. To estimate the core drive requirement, a useful formula is:

$$U \cong 100 + 750\left(\frac{\Delta B}{2.5}\right)^2\left(\frac{1\mu s}{t}\right)\left(\frac{d}{25\mu m}\right)$$

where U is the energy loss per cubic meter of core material, ΔB is the flux change, t is the pulse duration, and d is the thickness.

This loss formula is intended for flat voltage pulses, so the desired linear rise and fall would result in less loss than predicted for a 200 ns flat pulse, by about a factor of two. For a 200 ns pulse, the

expected loss is 3850J/m^3, in the metal. With the 64% assumed radial and axial packing factors, this becomes 2464J/m^3, and estimating the correction for the waveform as 0.5 gives 1232J/m^3. If each individual core is driven with 115 kV, and we desire an average of 1 MV/m, then 1.25 MV/m must be applied along the accelerating column to take into account the 20% of longitudinal space devoted to the focusing quadrupoles. This mean that a total of 84 voltage increments must be used, which may be all from axial or some combination of axial and radial subdivision. The total macroscopic core volume which is required is 0.23 m^3/m, so the total core loss would be 284 J/m and 3.4 J/voltage increment. The drive current that corresponds to the linear voltage rise does not have any simple analytical approximation; if we divide the 3.4 J losses by 200 ns, then the required power is 17 MW and the core current would be about 1 kA. To this must be added the beam current of 1.25 kA, for a total of 2.25 kA per voltage increment. From the assumed thyratron capabilities we would expect each pulser to drive about 3 or more of these voltage increments in parallel. These estimates are in the right vicinity as scaled from the Astron injector induction linac: this had a 300 ns pulse, a 14 kV voltage increment per core, and a current at the end of the pulse of 1.5 kA, for a flat voltage pulse. The inner diameter of those cores was very nearly half of the present one, so this would double the required current, on the other hand, the flat pulse versus ramped pulse correction halves the current, and the 200 ns versus 300 ns pulse duration difference increases the required current somewhat, in rough agreement with the 1 kA estimate above.

The intent here has been to show what an induction linac buncher might comprise, without any attempts at optimization. A further study involving the costs of the various choices would provide directions of where to concentrate effort. In particular, if it turns out that the pulsing system is where the costs are concentrated, then pulse compression schemes could be used to multiply the power from one pulser by a significant factor. If the costs are mostly associated with the length of the buncher, then more effort would be devoted to the bipolar pulsing possibility which could halve the length. The last observation is that many more beams could pass through any of the buncher linacs if desired.

List of Figures

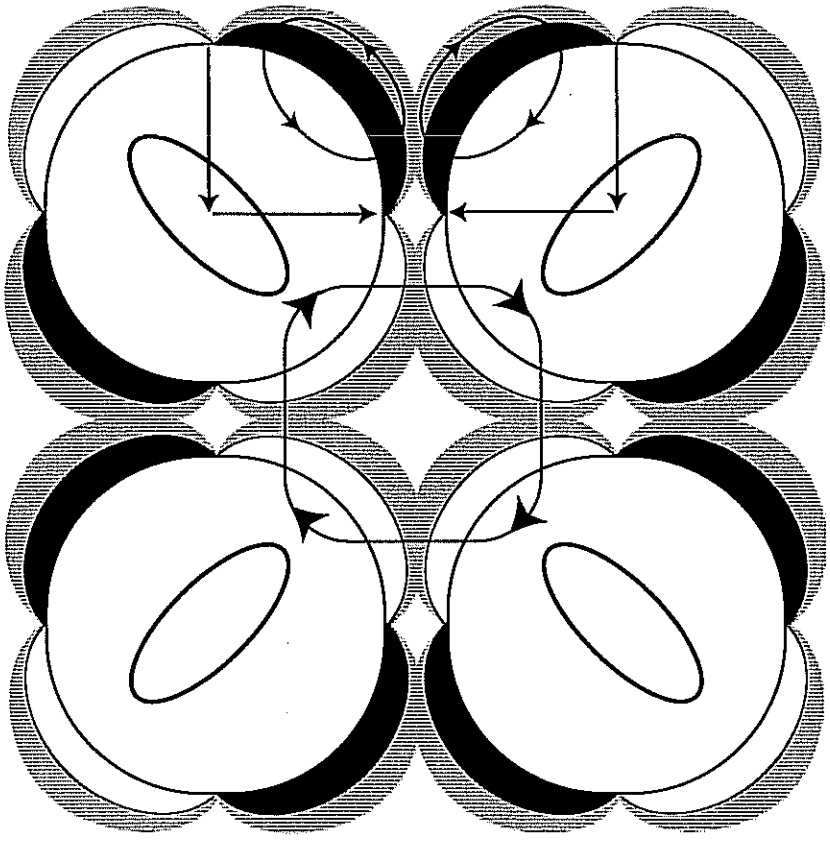

Figure 2.

Figure 1 A superconducting quadrupole for a single channel, showing the four coils and the compression collar. In an array the neighboring channels can be arranged such that the amount of superconductor and collar material can be reduced. A vacuum chamber and a layer of thermal insulation are placed within the superconducting coils.

Figure 2 The superconducting quadrupoles shown here occupy 20% of the half period length of 5 m. Most of the axial length is taken up by the accelerating column.

Figure 3 A 24-beam array, with space for field-terminating windings at the periphery of the arrary.

Figure 4 A typical induction module showing the core region and the drive geometry.

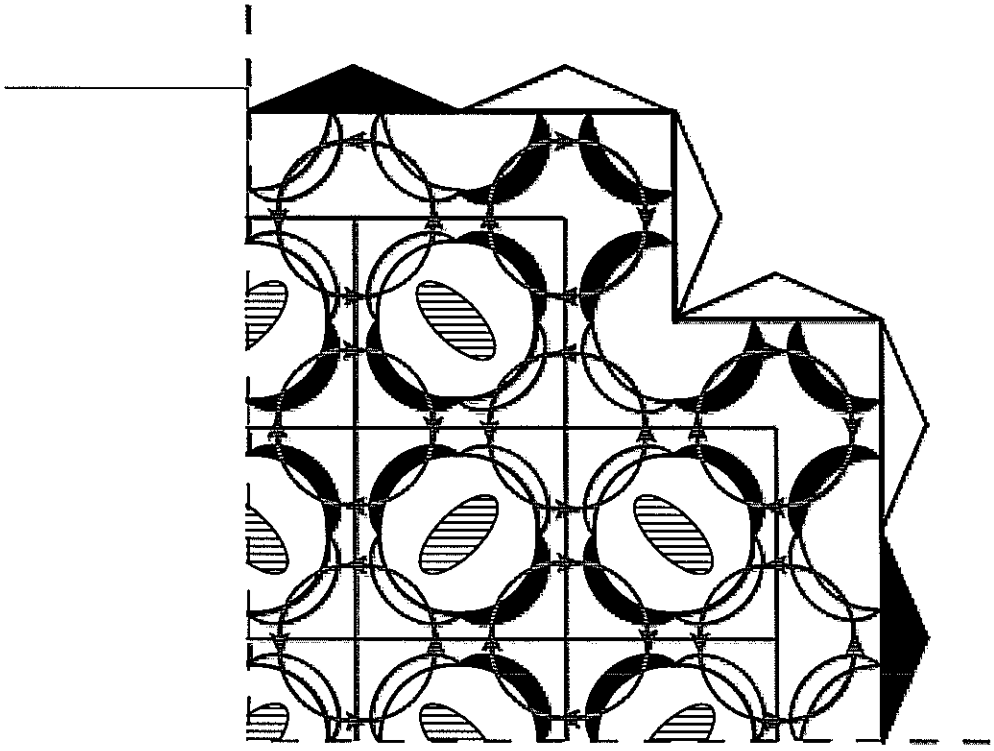

Figure 1

FIGURE 1

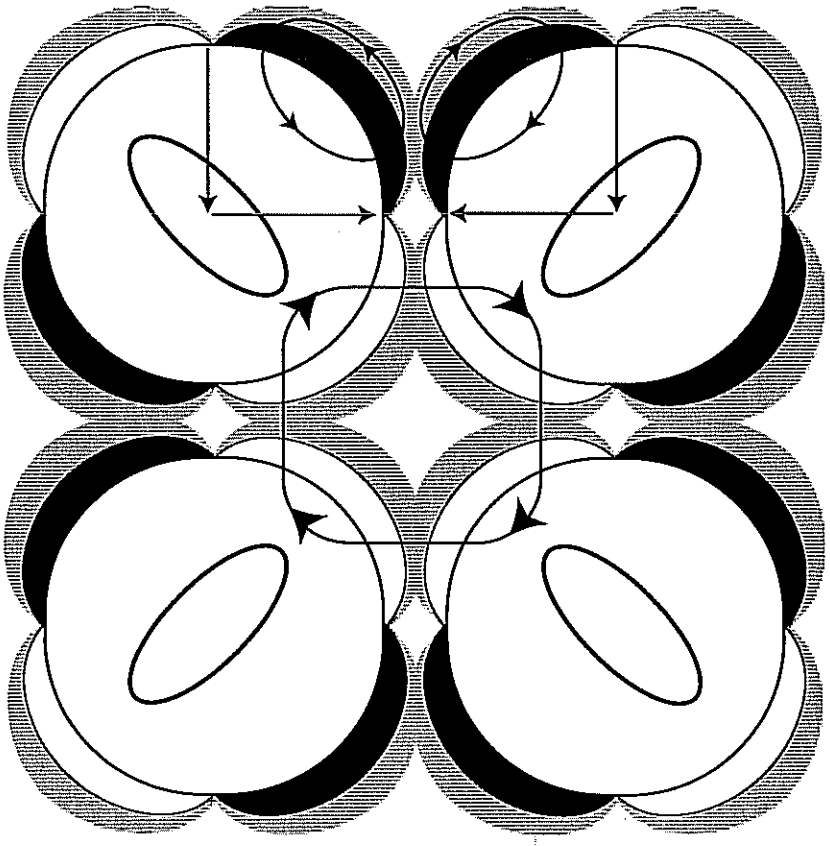

Figure 2.

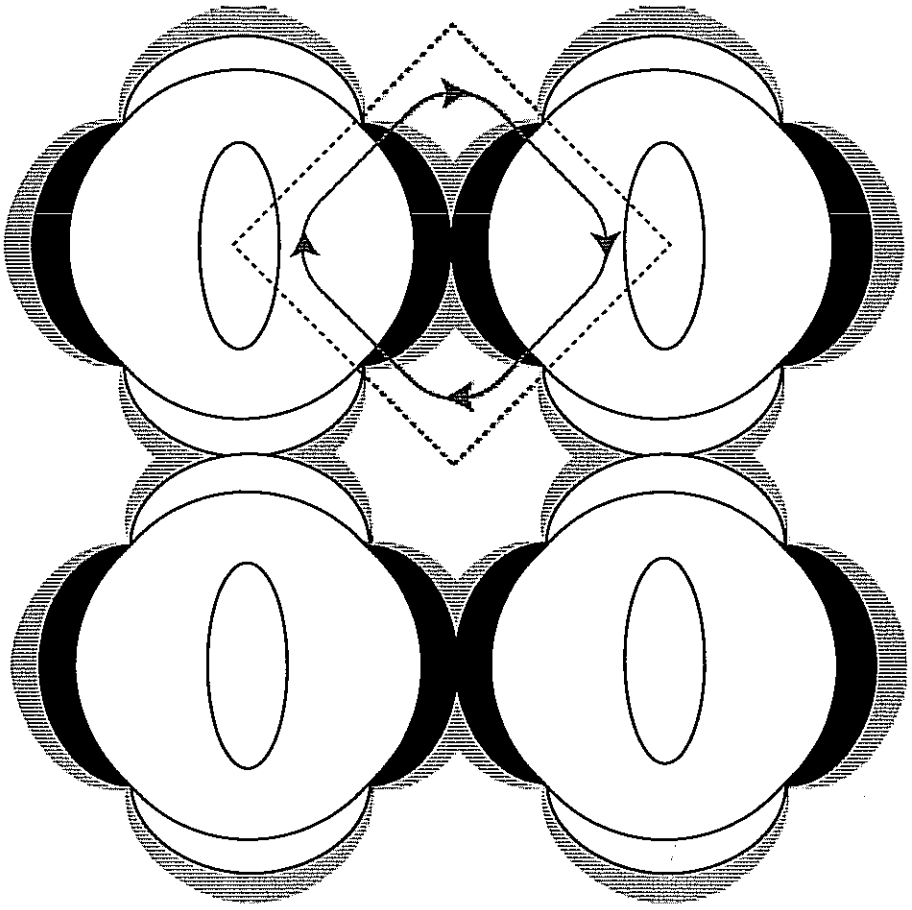

Figure 3