Proceedings of the 1963 Conference on Proton Linear Accelerators, New Haven, Connecticut, USA

October 20, 1963

INJECTION INTO THE BNL AGS

E. D. Courant
Brookhaven National Laboratory and Yale University

I wish to put down the over-all phase space and energy spread requirements that a linear accelerator must meet if it is to perform well as an injector for a synchrotron, and in particular, if it is to be a good 1 GeV injector for the existing AGS at Brookhaven.

The beam from a linac is characterized by its energy spread and its emittance. As we operate the AGS now, the duration of one turn of the beam is a little over 8 μsec, which is all we need. The energy spread must be less than that of the bucket for stable acceleration in the synchrotron. The maximum width of the momentum spread is

$$\frac{\Delta p}{p} = \left[\frac{eV \left[(\pi - 2\varphi_o) \sin \varphi_o - 2\cos \varphi_o \right]}{\pi h \beta^2 \gamma Mc^2 \left(\frac{1}{\gamma_t^2} - \frac{1}{\gamma^2} \right)} \right]^{1/2}$$

in which eV is the total voltage, h is the harmonic number, $1/\gamma_t^2$ is the momentum compaction factor, and φ_o is the stable phase angle in the synchrotron convention, where $\varphi_o = 0$ means no acceleration. Now numerically, if we use the same radio-frequencies that we have now at the AGS, the harmonic number is 12, and because the transition energy is large compared to

1 GeV,

$$\beta^2\gamma \left(\frac{1}{\gamma_s^2} - \frac{1}{\gamma^2} \right) = - \frac{\beta^2}{\gamma}$$

The function (β^2/γ) first increases with energy and then decreases again. The relative energy spread that is allowable at one GeV, compared to 50 MeV, is as follows: at 50 MeV, $\beta^2 = 0.1$ and $\gamma = 1$, while at 1 GeV, $\beta^2 = 0.75$ and $\gamma = 2$, so that the $\Delta p/p$ at 1 GeV is $(8/30)^{1/2}$ or about one-half of the 50 MeV value. With the values eV = 200 keV, $\varphi_o = 30^o$, h = 12, the allowable momentum spread at 1 GeV is $(\Delta p/p) = 3.3 \times 10^{-3}$ (half-width) so the full width is 6.6×10^{-3}. For efficient injection we want the energy spread to be somewhat smaller. In particular, there is a scheme of bunching by the synchrotron rf, which is used in both the CERN and the Brookhaven AGS accelerators, and which is, in principle, very closely related to the kind of matching we heard about from Dr. Teng. It requires an energy spread small compared to the above value. In this scheme, you inject the particles initially somewhere outside the bucket, and after a while they move together in phase space and then you clap the bucket around them. This sort of scheme gives high efficiency of capture -- as high as 70 or 80% -- but only if the energy spread is small compared to the width of the bucket. Therefore, I would say that a momentum spread of about 1×10^{-3} should be about the most that one should permit.

GLUCKSTERN: Are you assuming the same energy gain that you have now or twice that value?

COURANT: I am using the same energy gain that we presently have: V = 200 keV (applied), $\varphi_o = 30^o$, which means 100 keV/turn is required. There is a possibility which has been discussed at Brookhaven of having the magnetic field cycle speeded up by a factor of 2, which would imply applying twice as much voltage, in which case, $\Delta p/p$ would go up by a factor of $(2)^{1/2}$

Another possibility one might think of, would be to install a new rf system that operates at a higher harmonic, for example, at 200 Mc/sec, so that it would match directly to the injector. In that case you wouldn't have any bunching problem because the linac bunches would already be inside the buckets of the synchrotron. However, 200 Mc/sec is 50 times as high as the 4 Mc/sec that we use now, so that the harmonic number h would be 50 times 12, or 600, so that we would have to put a factor of $(50)^{-1/2}$ in $(\Delta p/p)$, reducing the total allowable bucket width to a maximum of 1×10^{-3} which would be imposing a much more severe requirement on the linear accelerator. The current thinking is very definitely in terms of leaving the rf system the same as at present or gradually changing from one to the other, but keeping the basic parameters the same.

Now consider emittance, and its relation to multiturn injection. One can inject as many turns as the ratio of the acceptance of the synchrotron to the emittance of the beam. In principle, one can do better than that, because one can have this ratio in horizontal phase space multiplied by the same ratio in vertical phase space and (possibly) in longitudinal phase space as well.

I will confine the discussion to radial phase space alone. First, let me indicate the requirement to be met. The whole purpose of the new injector is to enable us to get in more particles, and the reason why high-energy injection is good is because the space charge defocussing and similar phenomena become less serious as the particle energy goes up. Dr. Laslett has recently worked out a rather detailed theory of not only space charge, but image force defocussing, and finds that the space charge limit that the existing AGS can accept, with a reasonable beam size inside, is about 4 x 10^{13} particles/pulse. What we have been talking about, in our preliminary thinking for the injector, is a beam of the order of 1 to 2 x 10^{13}/pulse. That's about the best one can easily attain, assuming space charge defocussing only. We don't really know whether the space charge limit is the effective limit on the beam or whether there are other limits which are lower. Anyway we know that the linac can put out 50 mA because there exists at least one linac that does. The time for one turn at an injection energy of 1 GeV is about 3 μsec, so that 50 mA puts in about 10^{12} protons per turn. That means that in order to really fill the machine to the space charge limit we would need 40 turns.

Now, how do we get in 40 turns? Well, essentially we plan to use the scheme that we have just heard about from Dr. Curtis. The details are a little different because this is a different machine. We have a smaller aperture, but we're talking of higher energy. The emittance, E of the present 50 MeV linac, as Arie van Steenbergen has told you, is about 1.5 π cm-mrad at 50 MeV. At one GeV we assume (and have heard some evidence this morning that the assumption is more or less correct between 700 keV and 20 MeV, at least), that this emittance is inversely proportional to the particle momentum. In other words, if we can build a 50 MeV linac with the above emittance value, we should be able to build a 1 GeV linac at an emittance value of about 1/6th of the 50 MeV value, or 0.25 π cm-mrad. The acceptance of the synchrotron, a, can be written

$$a = \frac{\pi A^2}{\beta_{max}}$$

in which A is the semi-aperture, and β_{max} is the wave-length form factor. In our case, the useable horizontal semi-aperture, A, may be 2 in. At 1 GeV there is rather little orbit distortion caused by remanent field, so we have 5 cm of useable semi-aperture; β_{max} has a value of about 2200 cm, so that a $\simeq$ 11 π cm-mrad, and there seems to be room for 44 turns. I said we needed 40, so we've got 4 to spare -- there seems to be room, in principle, for the required 40 turns using horizontal phase space only.

I have not done as much detailed thinking on the exact mechanism of doing it as the MURA people have done, so I'm not going to do anything quite so elegant, but will present a very simple argument that indicates that one should be able to come fairly close to that ratio. The scheme is essentially what they had. Consider a circular orbit. It so happens that on either side of the AGS sector where we now do our injection, there are two straight sections that are just half a wavelength apart. If we put equal and opposite field bumps into those two straight sections, the orbit all the way around the machine will be unaffected, except for a little bump whose maximum is very nearly at the position of the injector, which is just ideal. This bump then can be made to be as much as 5 cm, and if we want 40 turns then the bump ought to shrink each turn by 1/40th of 5 cm which is 1.25 mm/turn. This turns out to require a rate of change of field that is rather easily attainable. One has to have a thin septum inflector, and to match the injected beam into the phase space of the synchrotron.

Now, the scheme I had thought about employs a resonance, and assumes a half-integral tune, rather than a quarter-integral. The machine was designed for a value of $\nu = 8\ 3/4$ but we can assume a value of $\nu \simeq 8\ 1/2$.

The width of the circulating beam is given by the formula

$$\Delta x \, \Delta x' = \frac{\epsilon}{\pi}$$

in which Δx is the beam width, $\Delta x'$ the angular spread, and ϵ the linac emittance. $\Delta x'$ is given by

$$\Delta x' = \frac{A}{\beta_{max}}$$

and Δx comes out to about 1 mm for the numbers we have. Now we can say that if we maintain a tune of 8 1/2, exactly (or exactly enough for 40 turns) these emittance ellipses are always tall in $\Delta x'$ and thin in Δx when they get back to the injector. Then we can forget about the height of the ellipse for a moment and draw the diagram shown in Fig. 1.

This shows the case of injection for 4, rather than 40 turns, for simplicity. The synchrotron radius, r, is shown horizontally, and the time, t, vertically, and so is the radial position of the septum. The equilibrium orbit as defined by the two field bumps is also shown. Suppose the width of each injected turn is such that it is twice the distance the equilibrium orbit moves every turn. Then, assume turn #0 has reversed itself, so half of it is where it gets chopped off by the septum, but turn #1 is right here to make up the loss. At time t_2, nothing is lost and #2 is injected.

The next time around (t_3), #1 is as shown; half of #0 and all of #2 are on the inside, and you're putting in #3. And so you see that this way we can fill the whole axis, and wind up with half of the last turn (here #4) lost, assuming the septum has zero thickness. If the septum has finite thickness, we lose the fraction of it which is equal to septum thickness divided by twice the change in orbit position per turn. Now our orbit setover per turn has to be about 1.5 mm and the septum thickness probably has to be 0.5mm, so this filling factor looks like about 80%.

Now let's look at the phase space as seen with respect to the actual equilibrium orbit, in Fig. 2. The alternate positions of the injected turns are shown, again for 4 turn injection. What about all the unoccupied space? Well, let's say we can concentrate the beam, so that each turn occupies a large x' spread up and down. We calculated 1 mm for Δx and have 50 mm in which to put these turns. Later turns, such as #3 and #4, will stick out past the synchrotron acceptance, and some will be lost altogether. The result of this effect will be an efficiency factor of $(\pi/4)$ (assuming an infinitesimal septum), because you can only fill a solid area with adjacent ellipses or circles with a maximum packing factor of $(\pi/4)$. This factor is to be multiplied by the septum efficiency of 80%. We can fill up 60 to 70% of the phase space in this way, which is probably as well as we need to do, assuming we have the 50 mA injector and a scheme like this one can indeed inject to maybe 3×10^{13} particles/pulse at 1 GeV, which is about the maximum number that we now think the AGS can accelerate.

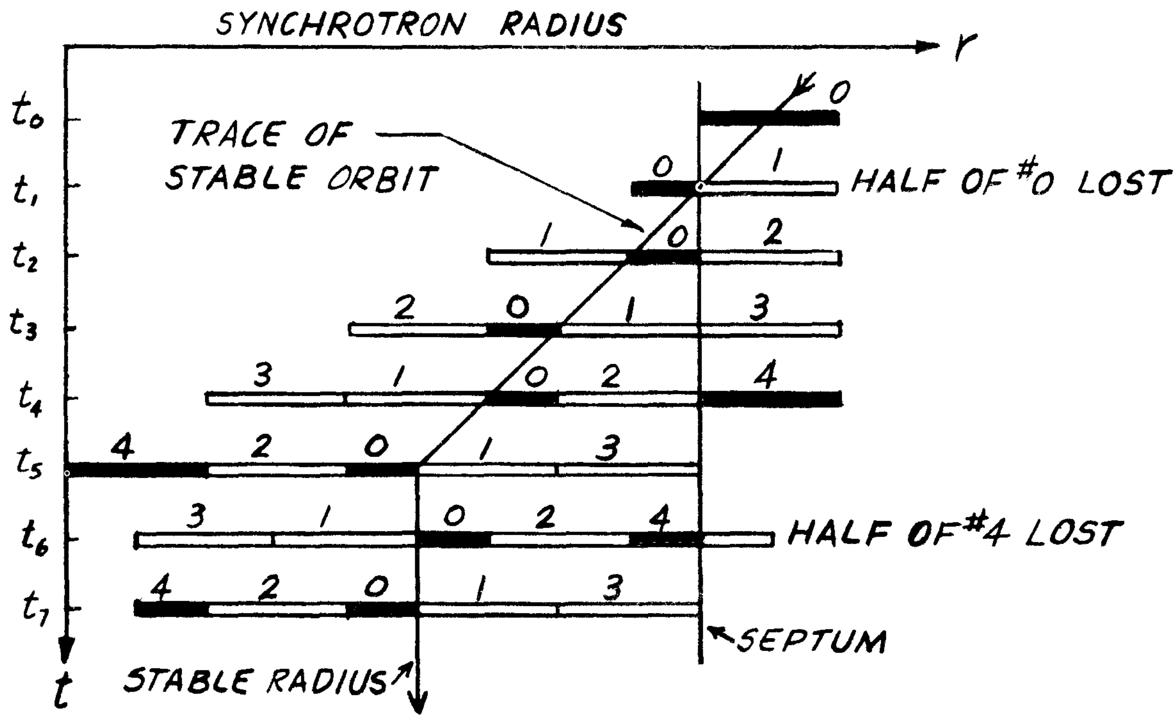

FIG. 1 - 4 TURN SYNCHROTRON INJECTION

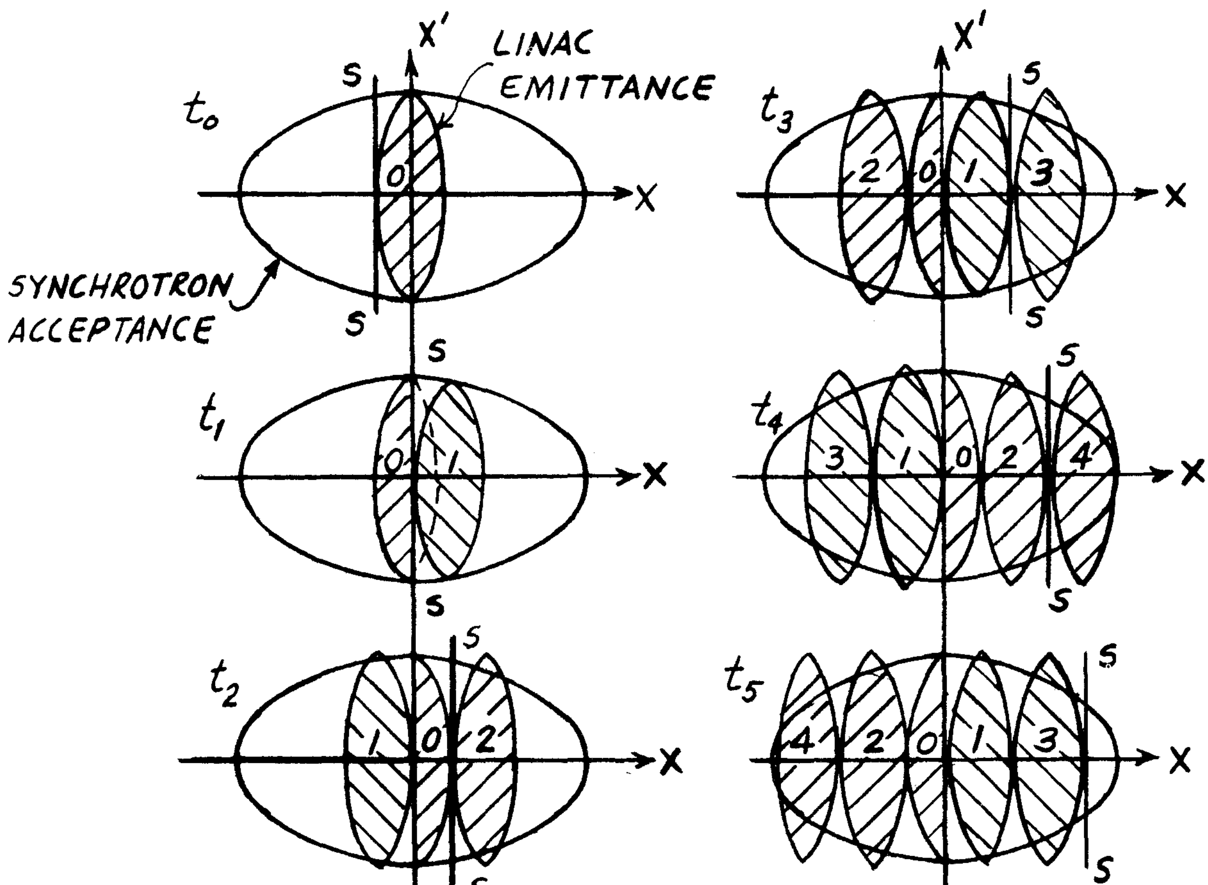

FIG. 2 - FILLING HORIZ. PHASE SPACE

In fact, one could do better than this, because one could, in principle, also introduce bumps in the field in the vertical direction and simultaneously make use of the same scheme vertically. The vertical phase-space ratio is not quite so favorable, but that means that for every turn we can inject in the horizontal plane, we could get in, let's say, three or four vertical turns.

BLEWETT: Do we really need a vertical bump to do that?

COURANT: You don't need a vertical bump to get two turns. But then you have a slightly hollow vertical phase space. You inject off the median plane, and everything has different phases of oscillations.

GLUCKSTERN: You can fill it up, probably, if you don't ask for this half integral tune, by population of different places.

COURANT: Well, I wanted to point out that filling up phase space can be done about as efficiently with a half-integral tune of the machine.

One thing to be watched is this. When you have a finite septum thickness and you use the scheme I suggest here, how serious a radiation hazard is it to have a large beam loss (of the order of 30 or 40%) at 1 GeV? Probably if one knows where it's going to be lost, one can judiciously place cleanup targets to be changed every so often.