and a decrease in the necessary rf voltage to maintain a given bunch intensity [43]. Further experiments are needed to fully demonstrate this technique.

The high rf voltage and beam power and the relatively small size of the machine require high gradient high power rf cavities. There is an agreement between Fermilab, BNL and KEK to collaborate on the research and development of such type of cavities. This work includes the study of magnet alloys and hybrid cavities using both ferrite and new magnet alloys, high power amplifiers and beam loading compensation.

The employment of barrier bucket rf cavities can effectively generate and manipulate a gap in the beam and reduce the space charge effect. Two barrier cavities have been built by BNL and KEK and will be installed on the AGS for test. Another high gradient barrier cavity using magnet alloys is under study at Fermilab.

IV. PION PRODUCTION, CAPTURE AND PHASE-ROTATION CHANNEL

This subsection first discusses the choice of target technology, optimization of the target geometry, and then describes design studies for the pion capture and phase-rotation channel.

Figure 4 gives an overview of the configuration for production of pions by a proton beam impinging on a long, thin target, followed by capture of low-momentum, forward pions in a channel of solenoid magnets with rf cavities to compress the bunch energy while letting the bunch length grow: thus rotating the bunch in phase space.

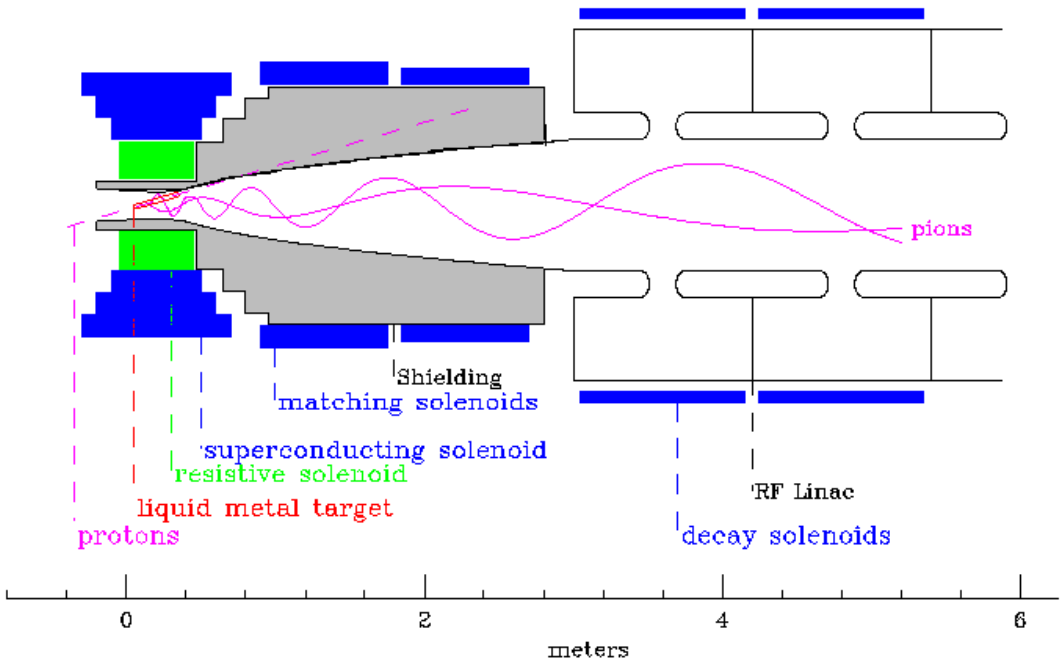

FIG. 4. Schematic view of pion production, capture and initial phase rotation: A pulse of 16-30 GeV protons is incident on a skewed target inside a high-field solenoid magnet followed by a decay and phase-rotation channel.

A. Pion production

To achieve the luminosities for muon colliders presented in Tb. I, 2×10^{12} (or 4×10^{12} in the 100 GeV CoM case) muons of each sign must be delivered to the collider ring in each pulse. We estimate that a muon has a probability of only $\frac{1}{4}$ of surviving the processes of cooling and acceleration, due to losses in beam apertures or by decay. Thus, 0.8×10^{13} muons (1.6×10^{13} at 100 GeV) must exit the phase rotation channel each pulse. For pulses of 2.5×10^{13} protons (5×10^{13} for 100 GeV), this requires 0.3 muons per initial proton. And since the efficiency of the phase-rotation channel is about 1/2, this is equivalent to a capture of about 0.6 pions per proton: a very high efficiency.

The pions are produced in the interaction of the proton beam with the primary target. Extensive simulations have been performed for pion production from 8-30 GeV proton beams on different target materials in a high-field solenoid

[11,44,45]. Three different Monte Carlo codes [46–49] predict similar pion yields despite significant differences in their physics models. Some members of the Collaboration are involved in an AGS experiment [50] to measure the yield of very low momentum pions, which will validate the codes in the critical kinematic region.

The pion yield is greater for relatively high-Z materials, and for these, the pion yield is maximal for longitudinal momenta of the same order as the average transverse momentum ($\approx$ 200 MeV/c). Targets of varying composition ($6 < Z < 82$), radii (0.4-3 cm) and thicknesses (0.5-3 nuclear interaction lengths (λ_I)) have been explored. For a fixed number of interaction lengths, the pion yield per proton rises almost linearly with proton energy [45], and hence almost proportional to the energy deposited in the target. The yield is higher for medium and high-Z target materials, with a noticeable gain at $Z > 26$ for 30-GeV proton beams, but with only a minor effect for $E \leq 16$ GeV, as shown in Fig. 5 where results of detailed MARS13(98) [47] simulations are presented. The curves show both the meson yield ($\pi + K$) from the targets in the 0.05≤p≤0.8 GeV/c momentum interval (labeled Y) and the number of mesons that are both captured in the high-field solenoid and transported into the decay channel (labeled YC). Typical statistical error is a few percent.

B. Target

The target should be 2-3 interaction lengths long to maximize pion production. A high-density material is favored to minimize the size and cost of the capture solenoid magnet. Target radii larger than about 1 cm lead to lower pion rates due to reabsorption, while smaller diameter targets reduce the added production from secondary interactions. Tilting the target by 100-150 mrad minimizes loss of pions by absorption in the target after one turn on their helical trajectory. Another advantage of the tilted target geometry is that the high-energy and neutral components of the shower can be absorbed in a water-cooled beam dump to the side of the focused beam (see Fig. 4).

About 30 kJ of energy is deposited in the target by each proton pulse (10% of the beam energy). Hence, the target absorbs 400 kW of power at the 15 Hz pulse rate. Cooling of the target via contact with a thermal bath would lead to unacceptable absorption of pions, and radiative cooling is inadequate for such high power in a compact target. Therefore, the target must move so as to carry the energy deposited by the proton beam to a heat exchanger outside the solenoid channel.

Both moving-solid-metal and flowing-liquid targets have been considered, with the latter as the currently preferred solution. A liquid is relatively easy to move, easy to cool, can be readily removed and replaced, and is the preferred target material for most spallation neutron sources under study. A liquid flowing in a pipe was considered, but calculations done at CERN [51] suggest serious problems in shock damage to the pipe. An open liquid jet is thus proposed.

A jet of liquid mercury has been demonstrated [51] but not exposed to a beam. For our application, safety and other considerations favor the use of a low-melting-point lead alloy rather than mercury. Gallium alloys, though with lower density, are also being considered. Experimental and theoretical studies are underway to determine the consequences of beam shock heating of the liquid. It is expected that the jet will disperse after being exposed to the beam. The target station must survive damage resulting from the violence in this dispersion. This consideration will determine the minimum beam, and thus jet, radius.

For a conducting liquid jet in a strong magnetic field, as proposed, strong eddy currents will be induced in the jet, causing reaction forces that may disrupt its flow [52,53]. The forces induced are proportional to the square of the jet radius, and set a maximum for this radius of order 5-10 mm. If this maximum is smaller than the minimum radius set by shock considerations, then multiple smaller beams and jets could be used; *e.g.*, four jets of 5 mm radius with four beams with 2.5 10^{13} protons per bunch. Alternatives include targets made from insulating materials such as liquid PtO_2 or Re_2O_3), slurries (*e.g.*, Pt in water), or powders [54].

A moving-solid-metal target is not the current baseline solution, but is a serious possibility. In this case [55] the target could consist of a long flat band or hoop of copper-nickel that moves along its length (as in a band saw). The band would be many meters in length, would be cooled by gas jets away from the target area, and would be supported and moved by rollers, as shown in Fig. 6.

The choice and parameters of the target are a critical question that needs resolution. Only an experiment in a magnetic field and in a beam will settle it, and this is being planned.

C. Capture

In order to capture all pions with transverse momenta $P_\perp$ less than their typical values of 200 MeV/c, the product of the capture solenoid field B and its radius R_a must be greater than 1.33 T-m. The use of a high field and small

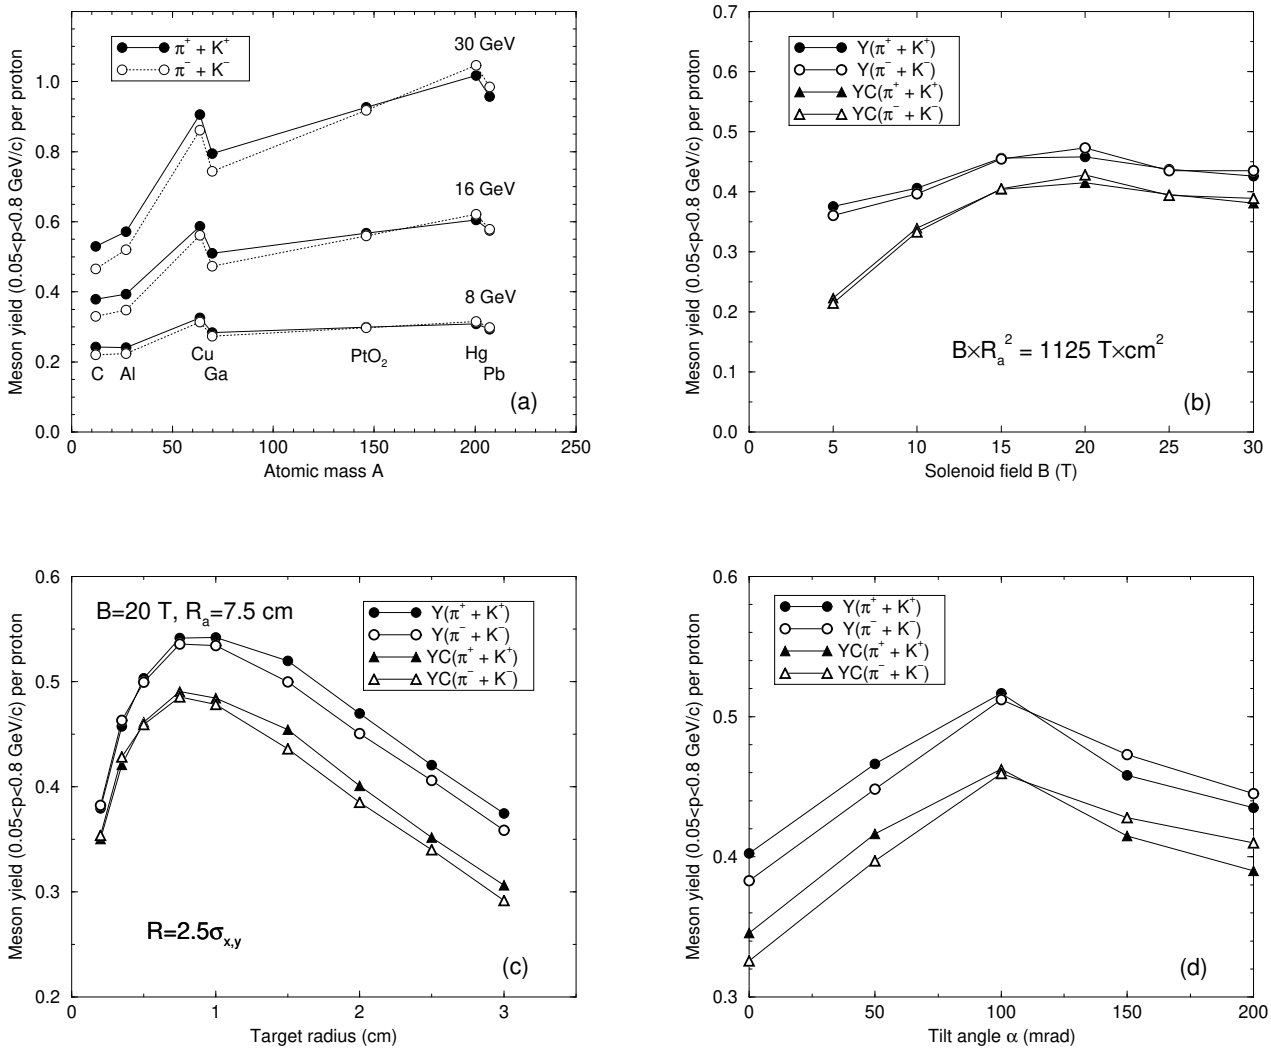

FIG. 5. Meson yield ($\pi + K$) from different targets tilted by angle α in a solenoidal field B of aperture R_a as calculated with the MARS13(98) code. The target is aligned along the beam. The curves labeled YC show mesons that are transported into the decay channel. (a) A 1.5-λ_I, 1-cm-radius target irradiated with 8, 16 and 30-GeV proton beams (σ_x=σ_y=4 mm) as a function of target atomic mass (B=20 T, R_a=7.5 cm, α=0); (b) A 3-λ_I, 1-cm-radius gallium target tilted at α=150 mrad in a 16-GeV proton beam (σ_x=σ_y=4 mm) *vs.* solenoid field for a fixed adiabatic invariant BR_a^2; (c) Yield as a function of a 3-λ_I gallium target radius in a 16-GeV proton beam (σ_x=σ_y=4 mm, B=20 T, R_a=7.5 cm, α=100); (d) Yield from a 3-λ_I, 1-cm-radius gallium target *vs.* tilt angle between the axis of the capture solenoid and the proton beam for a 16-GeV proton beam (σ_x=σ_y=4 mm, B=20 T, R_a=7.5 cm.

radius is preferred to minimize the corresponding transverse emittance, which is proportional to BR^2: for a fixed transverse momentum capture, this emittance is thus proportional to R. A field of 20 T and 7.5 cm radius was chosen on the basis of simulations described below. This gives BR = 1.5 T-m, BR^2 = 0.1125 T-m^2 and a maximum transverse momentum capture of $P_\perp = 225$ MeV/c.

A preliminary design [56] of the capture solenoid has an inner 6-T, 4-MW, water-cooled, hollow-conductor magnet with an inside diameter of 24 cm and an outside diameter of 60 cm. There is space for a 4-cm-thick, water-cooled, heavy-metal shield inside the coil. The outer superconducting magnet has three coils, with inside diameters of 60 to 80 cm. It generates an additional 14 T of field at the target and provides the required tapered field to match into the decay channel. Such an hybrid solenoid has parameters compatible with those of existing magnets [57].

The 20-T capture solenoid is matched via a transfer solenoid [44] into a decay channel consisting of a system of superconducting solenoids with the same adiabatic invariant $BR^2 \propto RP_\perp$. Thus, for a 1.25-T decay channel, B drops

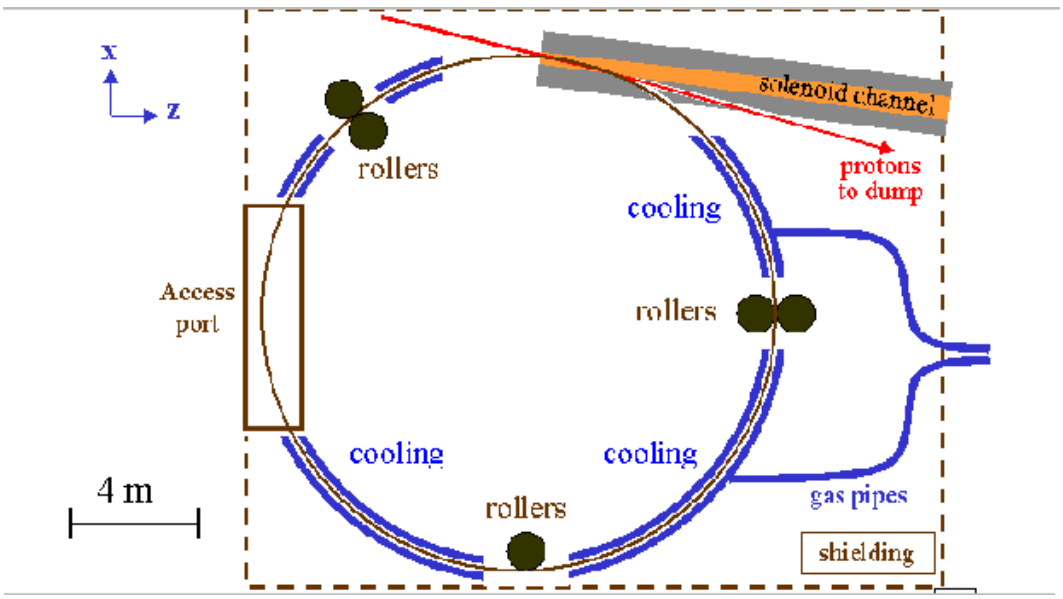

FIG. 6. Alternative concept of a solid metal target in the form of a rotating Cu-Ni band.

by a factor of 1/16 between the target and decay channel, R and $P_\perp$ change by factors of 4 and 1/4, respectively. This permits improved acceptance of transverse momentum within the decay channel, at the cost of an increased spread in longitudinal momentum. Figure 5(b) shows the meson yield as a function of field in the capture solenoid, with the radius of the capture solenoid adjusted to maintain the same BR_a^2 as in the decay channel. The optimum field is 20 T in the capture solenoid.

If the axis of the target is coincident with that of the solenoid field, then there is a relatively high probability that pions re-enter the target after one cycle on their helical trajectory and are lost due to nuclear interactions. When the target and proton beam are set at an angle of 100-150 mrad with respect to the field axis [45], the probability for such pion interactions at the target is reduced, and the overall production rate increased by 60 %, as shown in Fig. 5(d).

In sum, the simulations indicate that a 20 T solenoid of 16 cm inside diameter surrounding a tilted target will capture about half of all produced pions. With target efficiency included, about 0.6 pions per proton will enter the pion decay channel [45].

D. Phase Rotation Linac

The pions, and the muons into which they decay, have a momentum spread with an rms value of approximately 100 % and a peak value at about 200 MeV/c. It would be difficult to handle such a wide spread in any subsequent system. A linac is thus introduced along the decay channel, with frequencies and phases chosen to deaccelerate the fast particles and accelerate the slow ones; *i.e.*, to phase rotate the muon bunch. Several studies have been made of the design of this system, using differing ranges of rf frequency, delivering different final muon momenta, and differing final bunch lengths. In all cases, muon capture efficiencies of close to 0.3 muons per proton are obtained. Until the early stages of the ionization cooling have been designed, it is not yet possible to choose between them. Independent of the above choices is a question of the location of the focusing solenoid coils and rf cavity design. These will be discussed below in the subsection on Phase Rotation RF.

1. Lower Energy, longer bunch example

This example captures muons at a mean kinetic energy of 130 MeV. Table III gives parameters of the linacs used. It is seen that the frequencies vary between 60 and 30 MHz, and that the overall length is 42 m. Monte Carlo simulations [23], with the program MUONMC [58], were done using pion production calculated by ARC [46] for a copper target of 1 cm radius at an angle of 150 mrad. A uniform solenoidal field was assumed in the phase rotation, and the rf was approximated by a series of kicks.

TABLE III. Parameters of Phase Rotation Linacs

Linac	Length (m)	Frequency (MHz)	Gradient (MeV/m)
1	3	60	5
2	29	30	4
3	5	60	4
4	5	37	4

Figure 7 shows the energy *vs.* ct at the end of the decay and phase rotation channel. The abcissa is a measure of bunch length at the end of the channel: total transit time of each π/μ is multiplied by the velocity of light with total length of the channel (D) subtracted. Thus a photon at the center of the proton bunch at the target arrives at 0 cm. A loose final-bunch selection was defined with an energy 130 ± 70 MeV and bunch ct from 3 to 11 m. With this selection, the rms energy spread is 16.5 %, the rms ct is 1.7 m, and there are 0.385 muons per incident proton. A tighter selection with an energy 130 ±+ 35 MeV and bunch ct from 4 to 10 m gave an rms energy spread of 11.7 %, rms ct of 1.3 m, and contained 0.305 muons per incident proton.

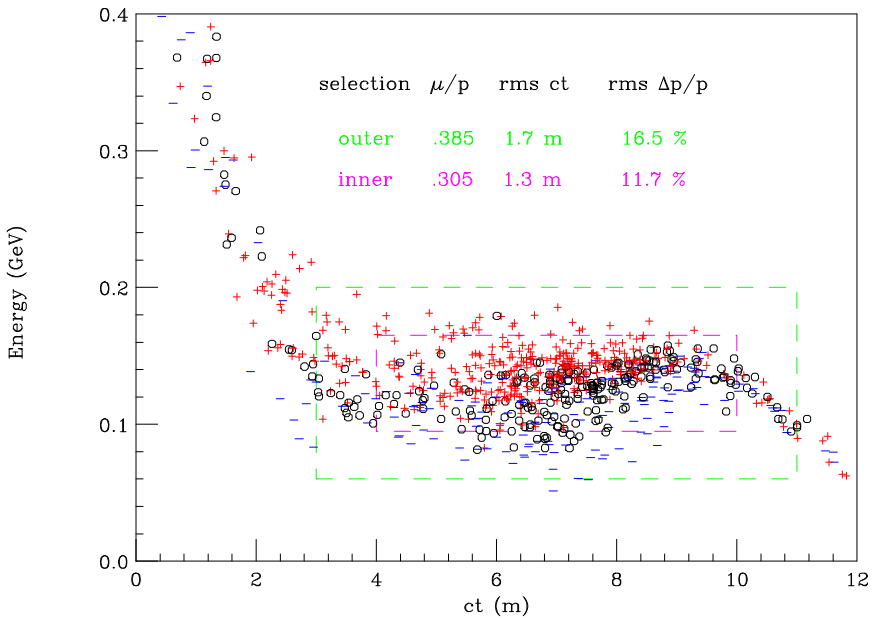

FIG. 7. Energy *vs.* ct of μ's at end of the Option-1 phase rotation channel. The symbols +, o and − denote muons with polarization $P > \frac{1}{3}$, $-\frac{1}{3} < P < \frac{1}{3}$ and $P < -\frac{1}{3}$, respectively.

2. Higher Energy, shorter bunch example

In this example the captured muons have a mean kinetic energy close to 320 MeV. It is taken from a Monte Carlo study [59] which uses the updated MARS pion production model [48] to generate pions created by 16 GeV protons on a 36 cm long, 1 cm radius co-axial gallium target. Figure 8 shows the longitudinal phase space of the muons at the end of an 80 m long channel with cavities of frequency in the 30-90 MHz range and acceleration gradients of 4-18 MeV/m. A total of 0.33 muons per proton fall within the indicated (6 m×300 MeV) cut. The rms bunch length inside the cut is 148 cm and rms energy spread is 62 MeV. Normalized 6-D emittance is 217 cm^3 and the transverse part is 1.86 cm.

A sample simulation with lithium hydride absorbers regularly spaced in the last 60 m of a 120 m decay channel and with compensating acceleration provided captures 0.297 muons with mean kinetic energy of about 380 MeV in a (6 m×300 MeV) cut. Longitudinal phase space is about the same as in the previous example but the transverse part shrinks to 0.95 cm due to ionization cooling which reduces 6-D phase space to 73.5 cm^3.

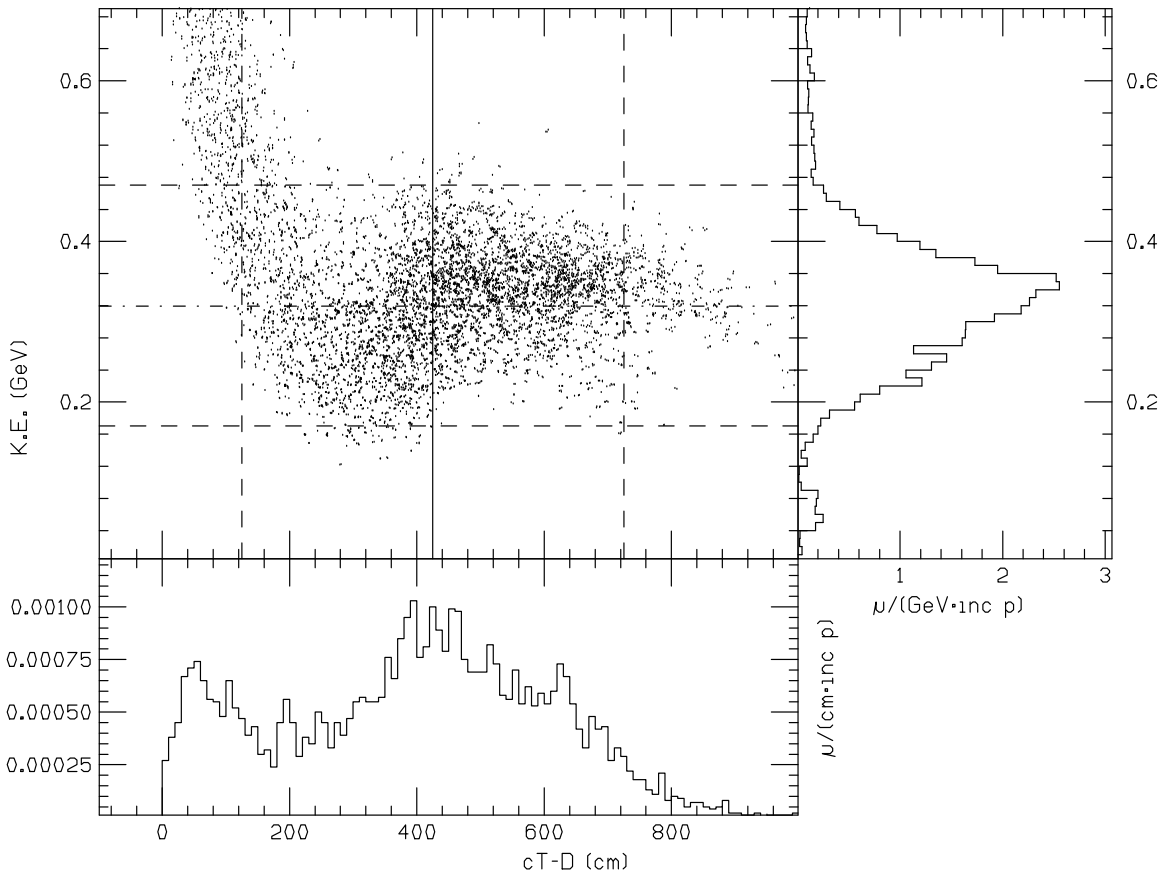

FIG. 8. Longitudinal phase space at the end of decay channel with projections onto time and energy axes. The four dashed lines delineate region deemed acceptable for the cooling channel.

E. Use of Both Signs

Protons on the target produce pions of both signs, and a solenoid will capture both, but the subsequent rf systems will have opposite effects on each sign. The proposed baseline approach uses two separate proton bunches to create separate positive and negative pion bunches and accepts the loss of half the pions/muons during phase rotation. If the pions can be charge separated with limited loss before the phase rotation cavities are reached, then higher luminosity may be obtained. The solenoid causes all charged particles to execute Larmor gyrations as they travel down the decay line. A gradient in the magnetic field or a curvature in the field produces drifts of the particle guide centers, and this can be used as the basis of a charge separation system. Drift directions are opposite for oppositely charged particles and perpendicular to the curvature or gradient. Drift velocities depend quadratically on particle velocity components.

The separation of charged pions in a curved solenoid decay line was studied in [44]. Because of the beam's resulting dispersion in a bent solenoid, an initial beam of radius r with maximum to minimum momentum ratio F will require a large beampipe of radius $(1+F)r$ downstream to accomodate the separated beams. A septum can then be used to capture the two beams into separate channels. Typically the reduction in yield for a curved solenoid compared to a straight solenoid is about 25% (due to the loss of very low and very high momentum pions to the walls or septum), but this must be weighed against the fact that both charge signs are captured for an overall net gain. A disadvantage is that this charge separation takes place over several meters of length during which time the beam spreads longitudinally. This makes capture in an rf phase rotation system difficult, although a large aperture cavity system could be incorporated in the bent solenoid region to alleviate this. The technique deserves further study and may be useful to consider as an intensity upgrade to a muon collection system.

F. Solenoids and RF

As noted above, capture using higher frequencies appears to be less efficient, and most studies now use frequencies down to 30 MHz. Such cavities, when conventionally designed, are very large (about 6.6 m diameter). In the Snowmass study [60] a reentrant design reduced this diameter to 2.52 m, but this is still large, and it was first assumed that the 5-T focusing solenoids would, for economic reasons, be placed within the irises of the cavities (see Fig. 9).

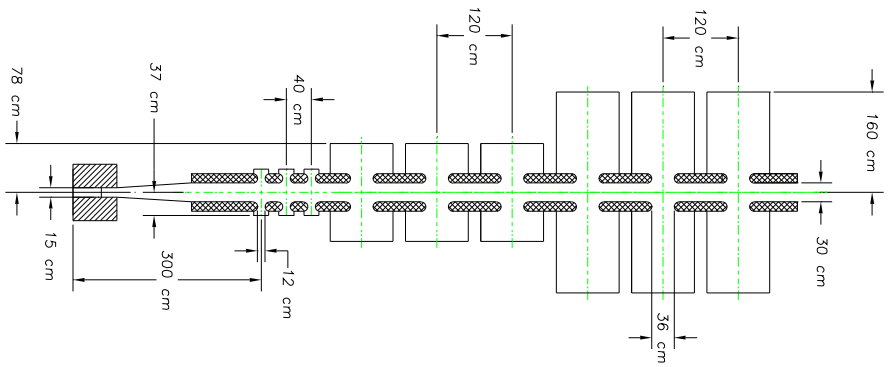

FIG. 9. Schematic of capture and phase rotation using rf cavities with superconducting solenoids (hatched) inside the irises. Only three sections each are shown for cavities operating at 90, 50, and 30 MHz

A study [61] of transmission down a realistic system of iris-located coils revealed betatron resonant excitation from the magnetic field periodicities, leading to significant particle loss. This was reduced by the use of more complicated coil shapes [60], smaller gaps, and shorter cavities, but remained a problem.

An alternative is to place continuous focusing coils outside the cavities (as shown in Fig. 4). In this case cost will be minimized with lower magnetic fields (1.25 - 2.5 T) and correspondingly larger decay channel radii (21 - 30 cm). Studies are underway to determine the optimal solution.

G. Polarization

Polarization of the muon beams presents a significant advantage over the unpolarized case, since signal and background usually come predominantly from different polarization states.

1. Polarized Muon Production

In the center of mass of a decaying pion, the outgoing muon is fully polarized ($P = -1$ for μ^+ and $+1$ for μ^-). In the lab system the polarization depends [62–64] on the decay angle θ_d and initial pion energy. For pion kinetic energy larger than the pion mass, the average polarization is about 20 %, and if nothing else is done, the polarization of the captured muons after the phase rotation system is approximately this value.

If higher polarization is required, some selection of muons from forward pion decays ($\cos\theta_d \rightarrow 1$) is required. Figure 7, above, showed the polarization of the phase rotated muons. The polarization $\{\mathrm{P} > \frac{1}{3}$, $-\frac{1}{3} < P < \frac{1}{3}$, and $\mathrm{P} < -\frac{1}{3}\}$ is marked by the symbols $+$, **o** and $-$ respectively. If a selection is made on the minimum energy of the muons, then greater polarization is obtained. The tighter the cut, the higher the polarization, but the less the fraction F_{loss} of muons that is selected. Figure 10 gives the results of a Monte Carlo study.

If this selection is made on both beams, and if the proton bunch intensity is maintained, then each muon bunch is reduced by the factor F_{loss} and the luminosity would fall by F_{loss}^2. But if, instead, proton bunches are merged so as to obtain half as many bunches with twice the intensity, then the muon bunch intensity is maintained and the luminosity (and repetition rate) falls only as F_{loss}.

One also notes that the luminosity could be maintained at the full polarized value if the proton source intensity could be increased. Such an increase in proton source intensity in the unpolarized case would be impractical because

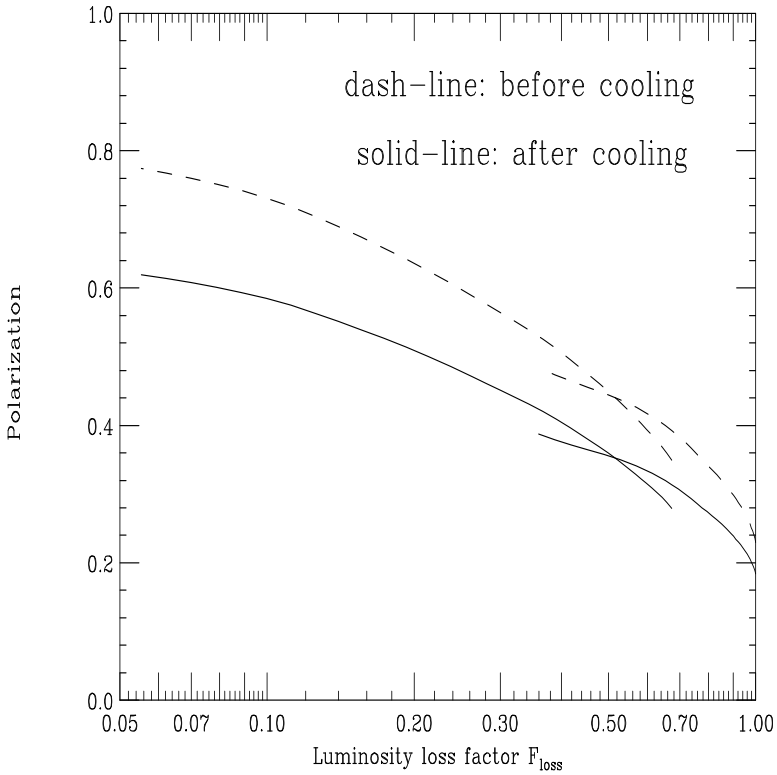

FIG. 10. Polarization *vs.* fraction F_{loss} of μ's accepted.

of the resultant excessive high energy muon beam power, but this restriction does not apply if the increase is used to offset losses in generating polarization.

2. Polarization Preservation

A paper [65] has discussed the preservation of muon polarization in some detail. During the ionization cooling process the muons lose energy in material and have a spin flip probability $\mathcal{P}$,

$$\mathcal{P} \approx \int \frac{m_e}{m_\mu} \beta_v^2 \frac{dE}{E}, \tag{2}$$

where β_v is the muon velocity divided by c, and $\frac{\Delta E}{E}$ is the fractional loss of energy due to ionization. In our case the integrated energy loss is approximately 3 GeV and the typical energy is 150 MeV, so the integrated spin-flip probability is close to 10 %. The change in polarization $\frac{\Delta \mathcal{P}}{\mathcal{P}}$ is twice the spin-flip probability, so the reduction in polarization is approximately 20 %. This loss is included in Fig. 10.

During circulation in any ring, the muon spin, if initially longitudinal, will precess by $\gamma(g-2)/2$ turns per revolution, where $(g-2)/2$ is 1.166×10^{-3}. A given energy spread $\frac{\Delta \gamma}{\gamma}$ will introduce variations in these precessions and cause dilution of the polarization. But if the particles remain in the ring for an exact integer number of synchrotron oscillations, then their individual average γ's will be the same and no dilution will occur.

In the collider, bending can be performed with the spin orientation in the vertical direction, and the spin rotated into the longitudinal direction only for the interaction region. The design of such spin rotators appears relatively straightforward, but long. This might be a preferred solution at high energies but is not practical in the 100-GeV machine. An alternative is to use such a small energy spread, as in the Higgs factory, that although the polarization vector precesses, the beam polarization does not become significantly diluted.

H. R&D Program

An R&D program is underway to continue theoretical studies (optimization of pion production and capturing up to the cooling channel) and to clarify several critical issues related to targetry and phase rotation. A jet of the room-temperature eutectic liquid alloy of Ga-Sn will be exposed to nanosecond pulses of 1.5×10^{13} 24-GeV protons at the F.E.B. U-line of the Brookhaven AGS to study the effect of the resulting pressure wave on the liquid. The same jet will also be used in conjunction with a 20-T, 20-cm-bore resistive magnet at the National High Magnetic Field Laboratory (Tallahassee, FL) to study the effect of eddy currents on jet propagation. Then, a pulsed, 20-T magnet will be added to the BNL test station to explore the full configuration of jet, magnet and pulsed proton beam. Also, a 70-MHz rf cavity will be exposed to the intense flux of secondary particles downstream of the target and 20-T magnet to determine viable operating parameters for the first phase-rotation cavity.

The first two studies should be accomplished during 1999, and the third and fourth in the years 2000/01.

V. IONIZATION COOLING

A. Introduction

The muon beam at the end of the decay channel is very intense, but diffuse in phase space. For a high luminosity collider, the phase-space volume of the muon beam must be reduced within a time of the order of the μ lifetime $\approx 2\ \mu s$. This requires the development of a new method for beam cooling.

Cooling by synchrotron radiation, conventional stochastic cooling and conventional electron cooling are all too slow. Optical stochastic cooling [66], electron cooling in a plasma discharge [67] and cooling in a crystal lattice [68,69] are being studied, but appear technologically difficult. Ionization cooling [70–72] of muons seems relatively straightforward, and it is unique for muons.

Ionization cooling involves passing the beam through some material in which the muons lose both transverse and longitudinal momentum by ionization loss (dE/dx). The longitudinal muon momentum is then restored by reacceleration, leaving a net loss of transverse momentum (transverse cooling). The process is repeated many times to achieve a large cooling factor.

The energy spread can also be reduced (longitudinal emittance exchange) by introducing a transverse variation in the absorber density or thickness (e.g. a wedge) at a location where there is dispersion (the position is energy dependent). This method results in an a corresponding increase of transverse momentum and is thus not true cooling, but an exchange of longitudinal and transverse emittances that, with transverse cooling, allows cooling in all dimensions.

The appropriate figure of merit for cooling is the final value of the 6-D relativistically invariant emittance ϵ_6 – the area in the 6-dimensional phase space (x, y, z, p_x, p_y, p_z)– since, to a fairly good approximation, it is preserved during acceleration and storage in the collider ring. Theoretical studies have shown that, assuming realistic parameters for the cooling hardware, ionization cooling can be expected to reduce the phase-space volume occupied by the initial muon beam by a factor of $10^5 - 10^6$. A complete cooling channel would consist of $20 - 30$ cooling stages, each stage yielding about a factor of two in phase-space reduction.

It is recognized that the feasibility of constructing a muon ionization cooling channel is on the critical path to understanding the viability of the whole muon collider concept. The muon cooling channel is the most novel part of a muon collider complex. Steady progress has been made both in improving the design of the channel and in adding detail to the computer simulations. A vigorous experimental program is needed to verify and benchmark the computer simulations.

In the following parts of this section we will briefly describe the physics underlying the process of ionization cooling, we will show results of simulations for some chosen examples and mention a six year R&D program to demonstrate the feasibility of using ionization cooling techniques to reduce the 6-dimensional phase space emittance.

B. Cooling Theory

In ionization cooling, the beam loses both transverse and longitudinal momentum as it passes through a material medium. At the same time its emittance is increased due to stochastic multiple scattering and Landau straggling. Subsequently, the longitudinal momentum can be restored by reacceleration, leaving a net loss of transverse momentum.

The approximate equation for transverse cooling (with energies in GeV) is