systematic error. Then, for example, from the decays of stored μ^-'s, the following neutrino oscillation channels could be studied by detection of the charged leptons from the interactions of neutrinos in the detector:

oscillation	detect
$\nu_\mu \to \nu_e$	e^-
$\nu_\mu \to \nu_\tau$	τ^-
$\bar{\nu}_e \to \bar{\nu}_\mu$	μ^+
$\bar{\nu}_e \to \bar{\nu}_\tau$	τ^+

The detected e^- or μ^+ have the "wrong sign" from the leptons produced by the interactions of the $\bar{\nu}_e$ and ν_μ flux. The known neutrino fluxes from muon decays could be used for long-baseline oscillation experiments at any detector on Earth. The probabilities for vacuum oscillations between two neutrino flavors are given by

$$P(\nu_a \to \nu_b) = \sin^2 2\theta \, \sin^2(1.27 \delta m^2 L/E) \tag{19}$$

with δm^2 in eV^2 and L/E in km/GeV. In a very long baseline experiment from Fermilab to the Gran Sasso laboratory or the Kamioka mine ($L = \mathcal{O}(10^4)$ km) with ν-energies $E_\nu = 20$ to 50 GeV ($L/E = 500$–200 km/GeV), neutrino charged-current interaction rates of $\sim 10^3$/year would result. In a long baseline experiment from Fermilab to the Soudan mine (L=732 km), the corresponding interaction rate is $\sim 10^4$/year. Such an experiment would have sensitivity to oscillations down to $\delta m^2 \sim 10^{-4} - 10^{-5}\,\mathrm{eV}^2$ for $\sin^2 2\theta = 1$ [58].

4. μp collider

The possibility of colliding 200-GeV muons with 1000-GeV protons at Fermilab is under study. This collider would reach a maximum $Q^2 \sim 8 \times 10^5$ GeV^2, which is $\sim$8 times the reach of the HERA ep collider, and deliver a luminosity $\sim 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, which is $\sim$300 times the HERA luminosity. The μp collider would produce $\sim 10^6$ neutral-current deep-inelastic-scattering events per year at $Q^2 > 5000$ GeV^2, which is more than a factor of 10^3 higher than at HERA. In the new physics realm, leptoquark couplings and contact interactions, if present, are likely to be larger for muons than for electrons. This μp collider would have sufficient sensitivity to probe leptoquarks up to a mass $M_{LQ} \sim 800$ GeV and contact interactions to a scale $\Lambda \sim 6$–9 TeV [89].

J. Summary of the physics potential

The First Muon Collider offers unique probes of supersymmetry (particularly s-channel Higgs boson resonances) and technicolor models (via s-channel production of techni-resonances), high-precision threshold measurements of W, t and SUSY particle masses, tests of SUSY radiative corrections that indirectly probe the existence of high-mass squarks, and a possible Z^0 factory for improved precision in polarization measurements and for B-physics studies of CP violation and mixing.

The Next Muon Collider guarantees access to heavy SUSY scalar particles and Z' states or to strong WW scattering if there are no Higgs bosons and no supersymmetry.

The Front End of a muon collider offers dramatic improvements in sensitivity for flavor-violating transitions (e.g., $\mu \to e\gamma$), access to high-Q^2 phenomena in deep-inelastic muon-proton and neutrino-proton interactions, and the ability to probe very small δm^2 via neutrino-oscillation studies in long-baseline experiments.

The muon collider would be crucial to unraveling the flavor dependence of any type of new physics that is found at the next generation of colliders.

Thus, muon colliders are robust options for probing new physics that may not be accessible at other colliders.

III. PROTON DRIVER

The overview of the required parameters is followed by a description of designs that have been studied in some detail. The section concludes with a discussion of the outstanding open issues.

A. Specifications

The proton driver requirements are determined by the design luminosity of the collider, and the efficiencies of muon collection, cooling, transport and acceleration. The baseline specification is for a 4-MW, 16-GeV or a 7-MW, 30-GeV proton driver, with a repetition rate of 15 Hz and 10^{14} protons per cycle in 2 bunches (for the 100-GeV machine) or 4 bunches (for the higher energies) of 5×10^{13} or 2.5×10^{13} protons, respectively. Half the bunches are used to make μ^- and the rest for μ^+ [90].

The total beam power is several MW, which is larger than that of existing synchrotrons. However, except for bunch length, these parameters are similar to those of Kaon factories [91] and spallation neutron sources [92]. As in those cases, the proton driver must have very low losses to permit inexpensive maintenance of components.

The rms bunch length for the protons on target has to be about 1 ns to: 1) reduce the initial longitudinal emittance of muons entering the cooling system, and 2) optimize the production of polarized muons. Although bunches of up to 6×10^{13} protons per cycle have been accelerated, the required peak current is 2000 A, which is unprecedented.

Since the collection of highly polarized μ's is inefficient (see section **IV.G**), the proton driver should eventually provide an additional factor of two or more in proton intensity to permit the luminosity to be maintained for polarized muon beams.

B. Possible options

Accelerator designs are site, and to some extent, time dependent, and there have been three studies at three different energies (30 GeV [93], 16 GeV [94] and 24 GeV [95–98]; see also [99]). In general, if the final energy is higher, the required currents are lower, bunch manipulation and apertures are easier, and the final momentum spread and space-charge tune shifts are less. Lowering the final energy gives somewhat more π's/Watt, a lower rf requirement ($V_{rf} \sim E^2$) and perhaps a lower facility cost.

In the low-energy muon collider, where two bunches of protons of 5×10^{13} are required on target, two bunches can be merged outside the driver. These two bunches would be extracted simultaneously from two different extraction ports, and fed by different transmission lines to the same target. By arranging the path lengths of the two lines appropriately, the two bunches can be exactly merged.

1. A generic design

A 7-MW collider-driver design based on parameters originally proposed in the Snowmass Feasibility study [93] consists of a 600-MeV linac, a 3.6-GeV booster and a 30-GeV driver. Both linac and booster are based on the BNL Spallation Neutron Source design [92], using a lower repetition rate and a lower total number of protons per pulse. For the 4-bunch case (2.5×10^{13} protons per bunch), the (95%) bunch area is assumed to be 2 eV-s at injection and < 4.5 eV-s at extraction. The driver lattice is derived from the lattice of the JHF driver using 90° FODO cells with missing dipoles in every third FODO cell, allowing a transition energy that is higher than the maximum energy or, perhaps, imaginary.

2. FNAL study

If a muon collider is built at an existing laboratory, then possibilities abound for symbiotic relationships with the other facilities and programs of that laboratory. For example, the proton driver for a muon collider might result from an upgrade of existing proton-source capabilities, and such an upgrade could then also enhance other future programs that use the proton beams.

Fermilab has conceived such a proton-source development plan [100] with three major components: an upgraded linac, a prebooster and a new booster, with the two boosters being rapid-cycling (15 Hz) synchrotrons. The two synchrotrons operate in series; the four proton bunches for the muon collider are formed in the prebooster and then accelerated sequentially in the prebooster and the booster. The plan could be implemented in stages, and other programs would benefit from each stage, but all three components are required to meet the luminosity goals of the muon colliders that have been considered so far.

Table III presents the major parameters of the two rings. Whenever the needs of the muon collider itself allow some flexibility, the parameters have been chosen to optimize the resulting facility as a proton source for the rest of

TABLE III. Baseline proton-driver parameters of the FNAL study.

	Linac	Booster	Driver
Energy range (GeV)	1	3	16
Rep. rate (Hz)	15	15	15
RF voltage per turn (MV)		0.15	1.5
Circumference (m)		158	474
Protons per bunch ($\times 10^{13}$)		2.5	2.5
Beam emittance [95%] (π mm-mrad)		200	240
Bunch area [95%] (eV-s)		1.5	< 2.0
Incoherent tune shift @ Inj.		0.39	0.39

the future program at Fermilab. For example, the machine circumferences and rf-harmonic numbers result in bunch trains that are compatible with the existing downstream proton machines.

A muon collider requires proton bunches that are both very intense and, at the pion-production target, very short. Strong transverse and longitudinal space-charge forces might disrupt such bunches in the synchrotrons unless measures to alleviate those effects are incorporated in the design. The Laslett incoherent-space-charge tune shift quantifies the severity of the transverse effects. A useful approximation for the space-charge tune shift $\Delta\nu_{sc}$ at the center of a round Gaussian beam is

$$\Delta\nu_{sc} = -\frac{3r_p N_{\rm tot}}{2\epsilon_n \beta\gamma^2 b} \tag{20}$$

In this expression $r_p = 1.535 \times 10^{-18}$ m is the so-called electromagnetic radius of the proton, $N_{\rm tot}$ is the total number of protons in the ring, ϵ_n is the 95% normalized transverse emittance, β and γ are the usual Lorentz kinematical factors, and $b \leq 1$ is the bunching factor, defined as the ratio of the average beam current to the peak current.

The approximation (20) implies that for a given total number of protons, here 10^{14}, the factors in the denominator are the only ways to reduce the tune shift to a specified maximum tolerable value, taken as 0.4. The bunching factor can be raised somewhat by careful tailoring of beam distributions, but here a typical value of 0.25 is conservatively assumed. Achieving the desired beam intensity then requires a combination of high injection energy, here taken as 1 GeV into the first ring, and large transverse normalized emittances, here assumed to be about 200π mm-mrad. The corresponding required aperture is about 13 cm in the first ring and about 10 cm in the second ring. With such large apertures in rapid-cycling synchrotrons, careful design of the beam pipes for both rings is required to manage eddy-current effects. Two approaches are under consideration. One is a thin Inconel pipe with water cooling and eddy-current coil corrections integrated on the pipe, as in the AGS Booster. The other is a ceramic beam pipe with a conductor inside to carry beam-image currents, as in ISIS.

The Fermilab linac presently delivers a 400-MeV beam, and is capable with modest modifications of accelerating as many as 3×10^{13} protons per cycle at 15 Hz [101]. A significant upgrade is required in order to deliver 10^{14} protons at 1 GeV. The energy can be raised by appending additional side-coupled modules to the downstream end of the linac. Increasing the linac beam intensity probably means increasing both the beam current and the duration of the beam pulse. Injection into the first ring is by charge stripping of the H^- beam; the incoming beam will be chopped and injected into pre-existing buckets to achieve high capture efficiency.

The circumference of the second ring is set equal to that of the existing Fermilab Booster. This choice provides several advantages. First, the new booster could occupy the same tunnel as a relocated Booster; secondly, the beam-batch length from a full second ring matches that of the present Booster, which simplifies matching to downstream machines for other programs. The output energy of 16 GeV then results from an assumed dipole packing fraction of 0.575 and a peak dipole field of 1.3 T, which is the highest dipole field that is consistent with straightforward, nonsaturating design of magnets having thin silicon-steel laminations. Driving such magnets into saturation would cause significant heating of the magnet yoke as well as potential problems with tracking between the dipoles and quadrupoles.

The prebooster also has 1.3-T dipole fields, and its circumference is one third that of the new booster; it operates at an rf harmonic number $h = 4$. The strategy for achieving the required short bunches at the target while alleviating space-charge effects in the rings is to start with four bunches occupying most of the circumference of the first small ring in order to keep the bunching factor large, and to do a bunch-shortening rotation in longitudinal phase space just before extraction from the second synchrotron. The four bunches are accelerated in the first ring to 3 GeV, then transferred bunch-to-bucket into the second ring with its harmonic number $h = 12$. At that energy, the kinematic factor in the tune-shift formula (20) is large enough to compensate for the smaller bunching factor in the second ring. The transfer energy of 3 GeV between the two rings roughly equalizes their space-charge tune shifts.

Both rings employ separated-function lattices with flexible momentum compaction in order to raise their transition energies above their respective extraction energies. This not only avoids having to accelerate beam through transition but also provides other advantages. Intense beams are not subject to certain instabilities such as the negative-mass instability below transition and empirically seem less susceptible to other instabilities such as the microwave instability. Also, the negative natural chromaticity is beneficial for stabilizing the beam below transition, thereby perhaps obviating the need for sextupole correctors, especially in the first ring. Having transition not too far above extraction also provides substantial bucket area in which to accomplish beam-shortening rf manipulations.

Several potential sources of instabilities in the rings have been examined [102]. Space charge is the main factor affecting the stability of the beams; the rings appear to be safe from longitudinal- and transverse-microwave instabilities. Of course, standard stabilizing methods such as active dampers are necessary to counteract some of the instabilities. Flexible momentum-compaction lattices would be useful not only to raise the transition energy above the extraction energy, but also to allow fast changes in the slip factor to facilitate bunch-narrowing manipulations at extraction time.

The magnet-power-supply circuit for each ring is a 15-Hz resonant system like that of the existing Booster, with dipoles and quadrupoles electrically in series. This implies that the second ring will accelerate only one batch at a time from the first ring, which is all that the muon collider needs. Adding about 15% of second harmonic to the magnet ramp reduces the required peak accelerating voltage by about 25%, which is probably worth doing, especially for the second ring with its large voltage requirement.

One of the advantages of a two-ring system is that the two rings divide the work of accelerating the beam. The rf system of the first ring is relatively modest because of its small circumference and small energy gain; that of the second ring is simplified because its high injection energy means a small rf-frequency swing [103].

ESME simulations of longitudinal motion show that the rms bunch length is 2 nsec as desired after the bunch rotation that occurs just before extraction from the second synchrotron.The bunch rotation creates momentum spreads of about 2% with longitudinal emittances of about 2 eV-s per bunch. Such spreads would contribute a few cm in quadrature to the beam size for a short period before extraction. This is thought to be tolerable, given the large apertures that are required in any case. High injection energies help to alleviate these longitudinal effects, which result from space-charge voltages having the same $1/\beta\gamma^2$ kinematic dependence as the transverse tune shifts.

3. AGS upgrade

The third study [95–98] is of an upgrade to the BNL AGS, which should produce bunches larger than those required for the muon collider, but at a lower repetition rate. The AGS presently produces 6×10^{13} protons in eight bunches at 25 GeV and 0.6 Hz. A 2.5-GeV accumulator ring in the AGS tunnel and AGS power-supply upgrade to 2.5-Hz operation would match the repetition rate to the 10-Hz repetition rate of the booster. This would generate 1 MW beam power. With an additional upgrade of the linac energy to 600 MeV, an intensity of 2×10^{14} protons per pulse in four bunches of 5×10^{13} at 25 GeV and 2.5 Hz could be reached, raising the power to 2 MW. The upgrades to the AGS accelerator complex are summarized in Fig.13. Other options are also under consideration, such as the addition of a second booster and 5-Hz operation, that would reach the baseline specification of 4 MW.

The AGS momentum acceptance of $\pm 3\%$ requires that the longitudinal phase space occupied by one bunch be less than 4.5 eV-s. This high bunch density in turn generates stringent demands on the earlier parts of the accelerator cycle. In particular, Landau damping from the beam momentum spread may guard against resistive wall instabilities during injection and longitudinal microwave instabilities after transition. Beam stability can be restored with a more-powerful transverse-damping system and possibly a new low-impedance vacuum chamber. The transverse microwave instability is predicted to occur after transition crossing unless damped by Landau damping from incoherent tune spread or possibly high-frequency quadrupoles.

C. Progress and open issues

Conventional rf manipulations appear able to produce 1- to 2-ns proton bunches if enough rf voltage to overcome the space-charge forces is used, and the beam energy is far enough from transition so the final bunch rotation is fast. Both simulations and experimental work have been directed at demonstrating that a short pulse can be produced easily.

An experiment at the AGS has shown that bunches with $\sigma_z = 2$ ns can be produced near transition from bunches with $\sigma_z \sim 8$ ns by bunch rotation [104,105]. In this experiment, the AGS was flattoped near transition (~ 7 GeV)

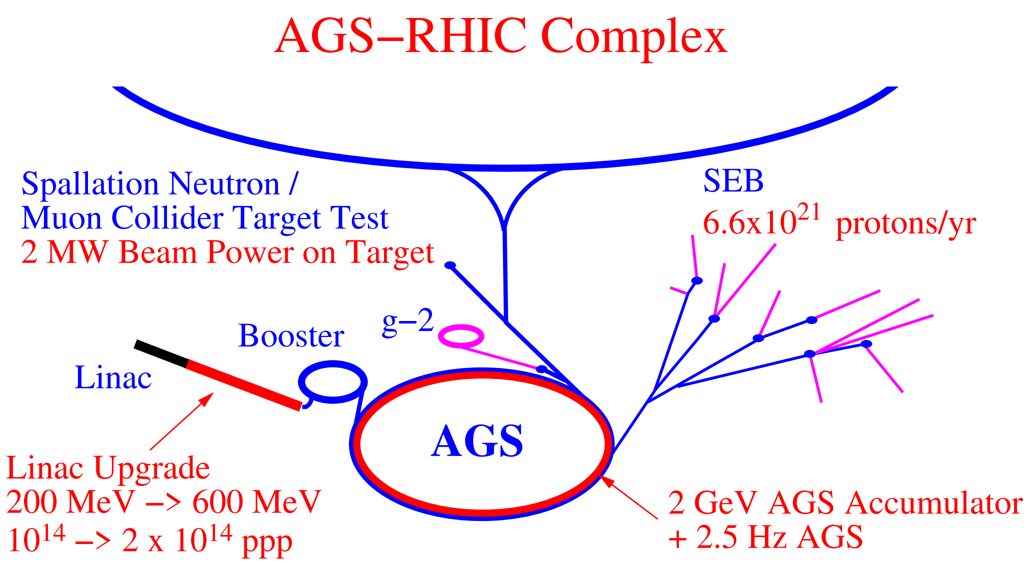

FIG. 13. The AGS–RHIC accelerator complex and a summary of possible intensity upgrades for the AGS.

while the γ_t-jump system was used to bring the transition energy suddenly to the beam energy, letting the bunch-energy spread expand and bunch length contract. The experiment also demonstrated that bunches are stable against microwave instabilities. In addition, the data were used to measure the lowest two orders of the momentum compaction factor.

The AGS bunch area, 1.5 eV-s, was comparable to that expected in the proton driver, but the bunch charge (though as large as 3-5 $\times$ 10^{12} protons) was only about one tenth of that required by the muon target. The proton driver would use a flexible momentum compaction lattice which would allow tuning far from transition and permitting a fast final bunch rotation [106]. In addition, the rf frequency would be higher than that of the AGS so the buckets (and bunches) would initially be only half as long. Thus bunch rotation could be expected to be easier with the new machine, which should compensate for the larger charge.

Simulations with the ESME code have also shown that 1-2 ns bunches of 5×10^{13} can be produced at extraction in a 16-GeV ring with the rf and emittance shown in Table III.

The efficiency of capturing and accelerating beam may be increased by compensation of the space-charge forces in the proton driver. The use of tunable inductive inserts in the ring vacuum chamber may permit active control and compensation of the longitudinal space charge below transition (since the inductive impedance is of the opposite sign to the capacitive space charge). Initial experiments at the KEK proton synchrotron and Los Alamos PSR [107] with short ferrite inserts appear to show a reduction in the synchrotron oscillation frequency shift caused by space charge and a decrease in the necessary rf voltage to maintain a given bunch intensity. Further experiments are needed to demonstrate this technique fully.

The high rf voltage and beam power and the relatively small size of the machine require high-gradient, high-power rf cavities. Fermilab, BNL and KEK are collaborating on 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 [108] rf cavities can effectively generate and manipulate a gap in the beam and reduce the space-charge effect. A successful test of this scheme has recently been completed [109], and two 40−kV barrier cavities have been built by BNL and KEK and are being installed on the AGS. Another high-gradient barrier cavity using magnet alloys is under study at Fermilab.